APRIL 1992

COMPUTER

WSI Bldg
$Y = \alpha x_1 + \beta x_2 + \gamma x_3 + \alpha x_4 + \beta x_5 + \gamma x_6$

WAFER-SCALE INTEGRATION:
Architectures and Algorithms

IEEE COMPUTER SOCIETY ®

THE INSTITUTE OF ELECTRICAL AND
ELECTRONICS ENGINEERS, INC. ®

NOTICE

WE'VE BEEN ASKED TO MAKE IT VERY CLEAR THAT
THERE'S A WHALE OF A DIFFERENCE BETWEEN QNX® AND UNIX®

WE'RE GLAD TO OBLIGE.

EVEN THOUGH THE QNX REALTIME OPERATING SYSTEM IS *NOT* BASED ON UNIX SOURCE CODE, WE CAN UNDERSTAND HOW SOME DEVELOPERS MIGHT CONFUSE THE TWO. AFTER ALL, QNX FOLLOWS THE LATEST IEEE POSIX 1003.1 AND 1003.2 OPEN SYSTEMS STANDARDS, SO YOU GET THE SAME API AND UTILITY SET FOUND IN MANY UNIX SYSTEMS.

BUT LOOK BENEATH THE SURFACE AND YOU'LL SEE TWO FUNDAMENTALLY DISTINCT ARCHITECTURES. A MONOLITHIC OS ON THE ONE HAND, A MICROKERNEL OS ON THE OTHER. THEY'RE DIFFERENT SPECIES ALTOGETHER, AS DIFFERENT AS A WHALE AND A SCHOOL OF DOLPHINS.

ARCHITECTURE YOU CAN BUILD ON

QNX'S MODULAR, MICROKERNEL ARCHITECTURE GIVES YOU REMARKABLE FLEXIBILITY. OS SERVICES ARE PROVIDED BY A TEAM OF COOPERATING PROCESSES RATHER THAN BY A MASSIVE MONOLITHIC KERNEL. THE RESULT? YOU CAN STRIP QNX DOWN TO A ROM-ABLE EMBEDDED SYSTEM. OR BUILD IT UP TO A VAST NETWORK AND HARNESS THE POWER OF HUNDREDS OF CPUS.

OUR PRICING IS MODULAR TOO, SO YOU'RE NOT FORCED TO PAY FOR WHAT YOU DON'T NEED.

REALTIME PERFORMANCE YOU CAN COUNT ON

NEED THE SPEED OF A FAST, REALTIME EXECUTIVE? QNX CLOCKS IN AT 16 μSEC PER CONTEXT SWITCH ON A 33 MHZ 80486. BUT QNX DOESN'T STOP THERE. WITH ITS PRIORITY-DRIVEN, PREEMPTIVE SCHEDULING, YOU CAN BUILD *REAL* REALTIME SOLUTIONS.

DISTRIBUTED PROCESSING YOU CAN BET ON

QNX CAN TRANSFORM A BUNCH OF ISOLATED MACHINES INTO A SEAMLESS SUPERCOMPUTER, ORCHESTRATING HUNDREDS OF CPUS WITH ITS NETWORK-WIDE IPC. YOU'RE IN COMPLETE CONTROL OF ALL RESOURCES AT ANY POINT, FROM THE PLANT FLOOR RIGHT UP TO YOUR DESKTOP.

QNX®

REALTIME OPERATING SYSTEM

SUPPORT YOU CAN DEPEND ON

DURING THE LAST TEN YEARS, WE'VE EARNED A REPUTATION FOR OUTSTANDING TECHNICAL SUPPORT. WE OFFER EVERYTHING FROM 24-HOUR ONLINE CONFERENCING TO ONSITE CONSULTING, SO YOU CAN EASILY REACH THE PEOPLE WHO MAKE UP THE QNX DEVELOPMENT TEAM. THEIR EXPERTISE CAN HELP YOU KEEP YOUR DEVELOPMENT PROJECTS RIGHT ON COURSE.

TO FIND OUT HOW YOUR APPLICATIONS CAN THRIVE IN THE QNX ENVIRONMENT, CALL 1-800-363-9001, (EXT. 108).

QUANTUM SOFTWARE SYSTEMS LTD.
175 TERRENCE MATTHEWS CRESCENT
KANATA, ONTARIO, CANADA
K2M 1W8
TEL: 613-591-0931
FAX: 613-591-3579

Reader Service Number 1

ARCHITECTURE MAKES THE DIFFERENCE

QNX IS A REGISTERED TRADEMARK OF QUANTUM SOFTWARE SYSTEMS LTD. UNIX IS A REGISTERED TRADEMARK OF UNIX SYSTEM LABORATORIES, INC. ©QUANTUM SOFTWARE SYSTEMS LTD. 1991

LANNET II.5
Local Area Networks

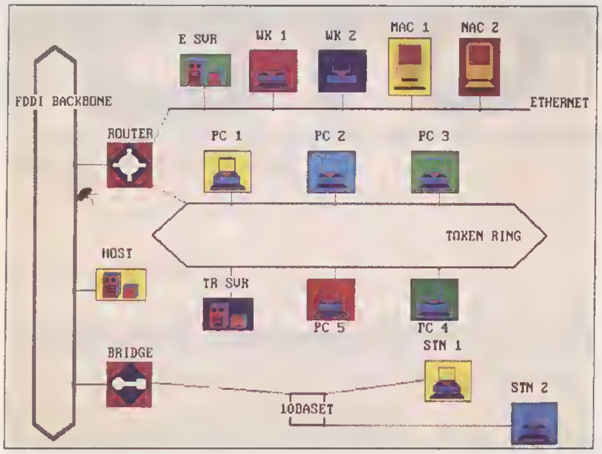

COMNET II.5
Wide Area Networks

New for local or wide area network analysts

Free trial and, if you act now, free training

LANNET II.5 uses simulation to predict your LAN performance. You simply describe your LAN and workload.

Animated simulation follows immediately --no programming.

Easy-to-understand results

You get an animated picture of your LAN. System bottlenecks and changing levels of utilization are apparent.

Your reports show LAN statistics such as transfer times, delays, and queues. Client, server, and gateway statistics show queue lengths, waiting times, and messages sent.

Your LAN simulated

You can predict the performance of any LAN. Industry standard protocols such as Ethernet, Token Ring, Token Bus, FDDI, and 10Base-T are built-in. Variations can be modeled.

COMNET II.5 uses simulation to predict your network performance. You simply describe your network, traffic load, and routing.

Animated simulation follows immediately --no programming.

Easy-to-understand results

You get an animated picture of your network. Routing choices and changing levels of network utilization are apparent.

Your reports show response times, blocking probabilities, call queueing and packet delays, network throughput, circuit group utilization, and circuit group queue statistics.

Your network simulated

You can include LAN's and multidrop lines in your model. X.25, SNA, DECnet, ISDN, SS7, fast packet, TCP/IP, token passing, and CSMA/CD are easily modeled.

Free Trial Offer

The free trial contains everything you need to try LANNET II.5™ or COMNET II.5® on your PC, Workstation, or Mainframe. Act now for free training**--no cost, no obligation.**

For immediate information

For LANNET II.5 call Eric Chapman, or for COMNET II.5 call Chris LeBaron, at (619) 457-9681, Fax (619) 457-1184. In Europe, call Nigel McNamara, in the UK, on 0276 671 671, Fax 0276 670 677. In Canada, call Peter Holt on (613) 782-2474, Fax (613) 782-2202.

University faculty should call about our special offer for research and teaching.

Mie. comp

Rush free trial and training information for:

☐LANNET II.5 ☐COMNET II.5

Name

Organization

Address

Telephone Fax

Computer Operating System

Return to:
CACI Products Co.
3344 N. Torrey Pines Ct.
La Jolla, CA 92037
Call Eric Chapman
or Chris LeBaron
at (619) 457-9681
Fax (619) 457-1184

In Europe:
CACI Products Division
Coliseum Business Ctr.
Watchmoor Park, Riverside Way
Camberley, Surrey GU15 3YL
United Kingdom
Call Nigel McNamara
on 0276 671 671
Fax 0276 670 677

In Canada:
CACI Products Div.
200-440 Laurier Ave. W.
Ottawa, Ontario K1R 7X6
Call Peter Holt
on (613) 782-2474
Fax (613) 782-2202

LANNET II.5 is a trademark and service mark of CACI Products Company. COMNET II.5 is a registered trademark and service mark of CACI Products Company. ©1992 CACI Products Company.

COMPUTER

April 1992 • Vol. 25, No. 4

ARTICLES

6 **Wafer-Scale Integration: Architectures and Algorithms**
Guest Editors' Introduction: W. Kent Fuchs and Earl E. Swartzlander Jr.

10 **Task-Flow Architecture for WSI Parallel Processing**
Robert W. Horst
This discussion explains the basics of task-flow architecture, examines a wafer-scale-integration implementation, and presents simulation results of a hypothetical machine.

19 **A Design and Yield Evaluation Technique for Wafer-Scale Memory**
Koichi Yamashita and Shohei Ikehara
Wafer-scale memory promises to provide greater yet cheaper storage volume than conventional memory systems using discrete chips. The authors used a Monte Carlo simulator for confirmation.

29 **Configuring a Wafer-Scale Two-Dimensional Array of Single-Bit Processors**
Ahmed Boubekeur, Jean-Luc Patry, Gabriele Saucier, and Jacques Trilhe
This overview of Europe's ELSA project shows how new hardware and software techniques can be used to achieve greater defect tolerance.

41 **Wafer-Scale Integration Using Restructurable VLSI**
Allan H. Anderson, Jack I. Raffel, and Peter W. Wyatt
Lasers are used in the final fabrication step for defect avoidance and customization. Nine wafer-scale systems illustrate the technique and its trade-offs.

50 **Wafer-Scale Transducer Arrays**
Glenn H. Chapman, M. Parameswaran, and Marek Syrzycki
The authors combined laser interconnection techniques and micromachining technology to produce wafer-scale transducer arrays. A prototype thermal-pixel dynamic scene simulator illustrates typical design requirements.

58 **Prospects for WSI: A Manufacturing Perspective**
Wojciech Maly
For 20 years, the drive toward smaller feature size has dominated the evolution of microelectronics. The cost analysis presented here suggests a change in this trend and better prospects for WSI.

66 **Project Overviews**
David L. Landis, Nitin Nigam, and Joseph W. Yoder; Vijay K. Jain, Hiroomi Hikawa, and D.C. Keezer
Overviews report progress in on-going research in two areas: the use of redundancy in WSI design and the development of a WSI architecture for rapid prototyping.

Published by the IEEE Computer Society

Circulation: *Computer* (ISSN 0018-9162) is published monthly by the IEEE Computer Society, 10662 Los Vaqueros Circle, PO Box 3014, Los Alamitos, CA 90720-1264; phone (714) 821-8380. IEEE Computer Society Headquarters, 1730 Massachusetts Ave. NW, Washington, DC 20036-1903; IEEE Headquarters, 345 East 47th St., New York, NY 10017. Annual subscription included in society member dues. Nonmember subscription rates: available upon request. Single copy prices: members $10.00; nonmembers $20.00. This magazine is also available in microfiche form.

Postmaster: Send undelivered copies and address changes to *Computer*, IEEE Service Center, 445 Hoes Lane, Piscataway, NJ 08855. Second class postage is paid at New York, New York, and at additional mailing offices. Canadian GST #125634188. Third class mail enclosed.

Bulk Rate
U.S. Postage
PAID
Shepherdsville, KY
Permit No. 93

Copyright and reprint permission: Copyright © 1992 by the Institute of Electrical and Electronics Engineers, Inc. All rights reserved. Abstracting is permitted with credit to the source. Libraries are permitted to photocopy beyond the limits of US copyright law for private use of patrons: (1) those post-1977 articles that carry a code at the bottom of the first page, provided the per-copy fee indicated in the code is paid through the Copyright Clearance Center, 29 Congress Street, Salem, MA 01970; (2) pre-1978 articles without fee. Instructors are permitted to photocopy isolated articles for noncommercial classroom use without fee. For other copying, reprint, or republication permission, write to Permissions Editor, *Computer*, 10662 Los Vaqueros Circle, PO Box 3014, Los Alamitos, CA 90720-1264.

Editorial: Unless otherwise stated, bylined articles, as well as product and service descriptions, reflect the author's or firm's opinion. Inclusion in *Computer* does not necessarily constitute endorsement by the IEEE or the Computer Society. All submissions are subject to editing for style, clarity, and space.

BPA

DEPARTMENTS

4 Letters to the Editor
January book review: A plea for McMathematics for engineers

76 Computer Society News
Development of ethical standards for computing professionals

79 Update

81 Product Reviews
PC networking

89 New Products

92 IC Announcements

93 Microsystem Announcements

94 Conferences

97 Call for Papers

100 Calendar

110 Book Reviews

112 Open Channel

Cover: Alexander Torres

Career Opportunities, 106; Computer Society Information, 88; Membership Application, 87; Change of Address Form, 4; Advertiser/Product Index, 80; Reader Service Card, 80A

In the next issue

Computer architectures for intelligent systems

COMPUTER
10662 Los Vaqueros Circle, PO Box 3014, Los Alamitos, CA 90720-1264

Editor-in-Chief
Jon T. Butler, Naval Postgraduate School
Associate Technical Editor
David C. Rine, George Mason University
Editorial Board
Dennis R. Allison, Stanford University
Laxmi N. Bhuyan, Texas A&M University
Thomas A. Corbi, IBM
Scott Davidson, AT&T
Richard Eckhouse, MOCO
Michael Evangelist, Andersen Consulting
Scott E. Fahlman, Carnegie Mellon University
John R. Gurd, University of Manchester
Brent T. Hailpern, IBM T.J. Watson Research Center
Tadao Ichikawa, Hiroshima University
Jogikal Jagadeesh, Ohio State University
Alan Kaminsky, Rochester Institute of Technology
Yale Patt, University of Michigan
Guylaine M. Pollock, Sandia National Laboratories
Norman F. Schneidewind, Naval Postgraduate School
Vincent Y. Shen, Hong Kong University of Science and Technology
Martha Sloan, Michigan Technological University
Will Tracz, IBM
Akinori Yonezawa, University of Tokyo

Submissions
Send 12 double-spaced copies of articles and special-issue proposals to Jon T. Butler, Dept. of Electrical and Computer Engineering, Code EC/Bu, Naval Postgraduate School, Monterey, CA 93943-5004, phone (408) 646-3299. For electronic submission, Butler's addresses are j.butler on Compmail II and butler@ece.nps.navy.mil on Internet.

Staff
Managing Editor: Marilyn Potes

Staff Editors: Christine Miller
Chuck Governale
Bob Carlson
Linda World
Contributing Editor: Ware Myers

Art Director: Jay Simpson
Design: Toni Van Buskirk
Production: Chris Patterson
Ken Duckworth

Publisher: True Seaborn
Editorial Director, CS Magazines: Marilyn Potes
Assistant to the Publisher: Pat Paulsen

Membership/Circulation: Christina Champion
Advertising Coordinators: Heidi Rex
Marian Tibayan

Magazine Operations Committee
James J. Farrell III (chair), Valdis Berzins, Jon T. Butler, B. Chandrasekaran, Carl Chang, Manuel d'Abreu, Dante Del Corso, John A.N. Lee, Peter R. Wilson

Publications Board
Harold Stone (chair), James J. Farrell III, Ronald G. Hoelzeman, Mary Jane Irwin, Ming T. (Mike) Liu, Michael C. Mulder, Theo Pavlidis, Sallie V. Sheppard, Kishor Trivedi

LETTERS TO THE EDITOR

January book review: A plea for McMathematics for engineers

To the editor:

Although Richard Hamming honors us by placing us in the august company of Dijkstra, Knuth, and Gries, he decidedly does not like our book (P.A. Fejer and D.A. Simovici, *Mathematical Foundations of Computer Science: Sets, Relations, Induction, Vol.I*, Springer-Verlag, New York, 1990). In his recent review in *Computer* (Jan. 1992, pp. 134-135), the single grudging trace of approval can be found near the end: Yes, the book might be useful, but only for those "mathematically inclined, logic-oriented" weaklings who, God forbid, aim to "publish papers"!

But in fairness, the reviewer's main fight is not with our book. He is embarked on a crusade against "purists" — that is, against people who believe, for example, that proofs by multiple trials are anything but proofs — and this review is simply another occasion to further his beliefs. Hamming believes that we favor proofs of program correctness. He is right in this respect. The other approach, which he clearly favors, is this: First you test the program, then you eliminate some mistakes, then you test the program again, then you eliminate the remaining mistakes (or, perhaps, the new mistakes introduced by the first round of fixing), then you test again, etc. When the program has no mistakes left, then the program is . . . correct!

While proving programs correct is difficult, the alternative has produced a vast amount of buggy software. What Hamming advocates is pragmatics not mathematics, arguments by hand-waving rather than mathematical rigor. In this context, it was gratifying to read his criticism of our Example 4.2.23, which is the *single* specific criticism in his review. Our example analyzes Quicksort as an application of strong induction. The algorithm makes use of a preexisting algorithm, Split, which is applicable to a nonempty sequence $\mathbf{a}$ of integers and yields an integer a and two sequences of integers $\mathbf{a}_R$ and $\mathbf{a}_L$ such that: (1) $\mathbf{a}_R a \mathbf{a}_L$ is a rearrangement of $\mathbf{a}$, (2) every entry in $\mathbf{a}_L$ is less than or equal to a, and (3) every entry in $\mathbf{a}_R$ is greater than or equal to a. And we continue:

> Given the input sequence $\mathbf{a}$, if $\mathbf{a}$ is the empty sequence, then $\mathbf{a}$ is already sorted, so the output for quicksort is the same as the input. Otherwise, apply the split algorithm to $\mathbf{a}$ to get $\mathbf{a}_L$, $\mathbf{a}_R$, and a. Recursively apply the quicksort algorithm to $\mathbf{a}_L$, $\mathbf{a}_R$ to produce outputs $\mathbf{b}_L$, $\mathbf{b}_R$, respectively. The output for quicksort applied to $\mathbf{a}$ is then $\mathbf{b}_L\langle a\rangle\mathbf{b}_R$.

The reviewer's criticism that the algorithm fails to deal correctly with the empty sequence is unfounded. As the first sentence above states, if $\mathbf{a}_R$ or $\mathbf{a}_L$ are empty, then quicksort returns that sequence.

Hamming is concerned about the sanity of a treatment where definitions suffocate the rest of the material. To make his point he uses a time-honored trick: citation out of context. He lists a sequence of definitions, theorems and corollaries, and "strictly mathematical examples" from a chapter that has a higher concentration of them than elsewhere in the book, then he complains about the number of them relative to the rest of the material. Numerically, the ratio of definitions to the rest of the material in this chapter is less than 13 percent, a number that speaks loudly of the thoroughness of Hamming's review. By selecting a fragment where definitions have a higher concentration, he conveys a falsehood.

Our book is not easy reading. It was written as a reference book, not as an introductory textbook, and it deals with fundamental and difficult topics (called by the reviewer "messy topics"). But it is a serious attempt to unify and systematize topics that are not easily found together, and it aims for clarity and attention to detail. On the other hand, we believe that the reviewer crossed the bounds of intellectual honesty and that reviews like his are a disservice to the public.

We hope that readers will recognize Hamming's writing for what it is: a plea for bargain-basement McMathematics for engineers and a lowering of intellectual standards in mathematical education.

Dan A. Simovici
University of Massachusetts at Boston

Reviewer reply:

On reading Simovici's "Plea," I am reminded of the story of the man who observed to his friend that the friend's mistress had a mole on her cheek. He was immediately accused of hating the lady. Similarly, to me at least, the "Plea" has little relationship to what I wrote and accuses me of opinions I do not have. I leave interested readers to reread my review and ask themselves how well the above story fits the situation. Those who are curious about some details can read the book itself and make up their own minds.

Richard W. Hamming
Naval Postgraduate School
Monterey, Calif.

Moving?

Name (Please Print)

New Address

City State/Country Zip

PLEASE NOTIFY US 4 WEEKS IN ADVANCE

MAIL TO:
IEEE Service Center
445 Hoes Lane
Piscataway, NJ 08854

ATTACH LABEL HERE

- This notice of address change will apply to all IEEE publications to which you subscribe.
- List new address above.
- If you have a question about your subscription, place label here and clip this form to your letter.

If you don't think our real-time systems have taken off, consider some of the programs we've landed.

You already know the significance of the Alsys name when it comes to quality Ada. But you may not have realized that over half of our business is in—and over half of all our resources are devoted to—real-time applications. Right now Boeing, Lockheed, McDonnell Douglas, NASA, AIRBUS, European Space Agency and many others rely on us for solutions for mission and safety-critical applications. Alsys offers one of the broadest ranges of compiler products and developer tools, plus unparalleled expertise and guidance every step of the way.

If you need to reduce risk, deliver projects on time, and do both within budget, the first step is easy: send in the coupon, or call 617-270-0030. We'll rush you a handbook detailing some powerful reasons why you should consider enlisting our support. Alsys. We're serious about real-time.

Please rush me the Alsys handbook detailing your real-time capabilities.

(name) (title)

(company) (phone no.)

(address)

Alsys, Inc., 67 South Bedford Street, Burlington, MA 01803-5152,
phone: 617-270-0030 fax: 617-270-6882

Alsys maintains major operations
in the U.S., Sweden, France,
England, Germany and Japan.

alsys
A Thomson—CSF company

IE

Reader Service Number 2

Guest Editors' Introduction

Wafer-Scale Integration: Architectures and Algorithms

W. Kent Fuchs
University of Illinois

Earl E. Swartzlander Jr.
University of Texas at Austin

Growth in the physical area of integrated circuits is the next great frontier in implementing complex microelectronic systems. Wafer-scale integration (WSI) pushes the physical dimensions of integrated circuits to their natural limit, the size of a wafer. There are many definitions of wafer-scale integration, ranging from chips large enough to require fault circumvention in order to achieve an acceptable yield (for example, state-of-the-art DRAMs) to chips that are a full slice of a wafer (for example, the first monolithic logic gates constructed on half-inch wafers).

Traditional very large scale integration (VLSI) circuits are created by fabricating a wafer, testing the individual die, dicing the wafer, and packaging the defect-free chips. A microelectronic system is implemented by mounting the individually packaged chips on printed circuit boards. By contrast, in monolithic WSI a wafer may be fabricated with several types of circuits; the circuits are tested and the defect-free circuits are interconnected to realize the system on the wafer.

Since the invention of the integrated circuit over 30 years ago, research and development has focused on increasing the number of transistors contained in a single IC.[1] The dramatically increasing performance of ICs in the last two decades has been due to both the growth in the number of transistors in each IC and the enhanced speed of transistors and interconnect. The increase in the number of transistors has been due primarily to reduced feature sizes and larger IC area. The architectures and algorithms used to exploit and enhance this growing physical area of ICs are the focus of this special issue.

0018-9162/92/0400-0006$03.00 © 1992 IEEE

Successful application of wafer-scale integration

Wafer-scale integration is not a panacea for the difficulties and cost in VLSI fabrication or for the growing demand for greater system performance. There are several requirements for the application of WSI to successfully compete with alternative technologies:[5]

(1) WSI fabrication must use state-of-the-art processing with the smallest possible feature sizes.

(2) Complexity in processing beyond that required for competing VLSI and alternative packaging technologies should be avoided.

(3) Architectures and applications must be crafted to avoid long interconnect delays and to exploit nearest neighbor connections.

(4) Methods for accurately testing, diagnosing, and repairing failures must be embedded in the design and fabrication process.

(5) The product must include the latest advances in systems architectures and applications.

WSI advantages and challenges

There is significant benefit to increasing the number of transistors in a single IC.[1-6] These include

1. higher system clock rates due to fewer off-chip connections and shorter intermodule connections,
2. lower power requirements due to fewer output pad drivers,
3. enhanced system reliability due to fewer interpackage connections, reduced bonding requirements, and fewer packages, and
4. reduced volume and weight.

The primary challenges to WSI include market economics as well as technology hurdles. The existing investment in competing VLSI products and alternative packaging technologies could impede the economic success of WSI products.[7] The primary technology hurdles include power dissipation, the need for reasonable manufacturing yield, difficulties with mask lithography fabrication, testing requirements, and potential signal delays due to long on-wafer interconnect. The articles in this special issue describe approaches to meeting these challenges.

WSI successes and failures

Since the creation of the first IC in the late 1950s, research has been conducted on whole-wafer circuits.[8] In the early 1960s, companies such as Texas Instruments, Hughes, and Westinghouse worked to enlarge the physical area of the IC and thereby increase the number of transistors.[1-6] Several companies produced WSI products during the last two decades, including memories, solid-state disk replacements, processor arrays, and neural networks.[2,4,5] Some dropped their WSI product lines because of combined difficulties with technology and the competition from VLSI. One well-known example of a commercially unsuccessful WSI product is the IBM-compatible computer by Trilogy.[3,6]

Some companies discontinued products because they failed to maintain state-of-the-art processing and packaging. WSI implementations include methods for enhancing yield through design for defect tolerance. WSI also increases the number of transistors by increasing the size of the IC. However, WSI is not a substitute for state-of-the-art processing with the smallest possible feature sizes and fewest possible defects.

Converging technologies

Microelectronic system packaging technologies are converging. Increasingly popular multichip modules (MCMs) can be thought of as hybrid or flip-chip WSI.[5,7] In hybrid WSI and MCMs, good dies are densely mounted on an interconnect substrate. Hybrid WSI has traditionally used a silicon substrate for interconnection, while MCM technology has used a variety of materials for packaging and interconnect. MCMs and hybrid WSI can also be viewed as a natural extension of aggressive printed circuit board technology.

Many of the difficulties faced by MCMs and hybrid WSI are similar to those faced by monolithic WSI.[7] Problems with heat dissipation, power and clock distribution, and system design are all common to MCMs and WSI. Current MCMs do not use the defect tolerance techniques necessary in monolithic WSI for achieving acceptable yield. However, because of their high cost, MCMs do require fault location and repair after burn-in of the package. It is likely that MCMs will use increasingly larger chips, incorporating defect tolerance, and for some applications defect tolerance will be incorporated into the MCM design.

All forms of IC technology in the next 25 years will benefit from the research and development in the design, fabrication, testing, and packaging of WSI that occurred over the past 25 years. The need for more transistors will probably cause ICs to grow dramatically in area. Design and implementation techniques for solving problems now considered unique to WSI, such as yield, packaging, heat dissipation, and interconnect delay, will be common to all IC implementations. The articles in this special issue describe breakthroughs in WSI technology, architectures, and algorithms that will affect future generations of integrated microelectronic systems. ■

Acknowledgments

We thank Vicki McDaniel and Shiree Schade for their assistance in managing the article submissions, reviews, and mailings for this special issue. We also thank Jon Butler for making the issue possible, and for his guidance as editor-in-chief. Carla Shaw made valuable contributions by helping us manage the schedule. Finally, we are sincerely grateful for the hard-working anonymous referees and the authors of all the submitted papers.

References

1. W.R. Moore, "A Critical Review of Fault-Tolerant Chips and WSI," in *Wafer-Scale Integration*, C. Jesshope and W. Moore, eds., Adam Hilger, Bristol, England, 1986.

2. R.M. Lea, "Wafer-Scale Integration: Motivation, Perspective, and Potential," *Proc. IFIP WG 10.5 Workshop on Wafer-Scale Integration, II*, Elsevier North-Holland, New York, 1988, pp. 3-17.

3. D.L. Peltzer, "Wafer-Scale Integration: The Limits of VLSI," *VLSI Design*, Vol. 4, No. 5, Sept. 1983, pp. 43-47.

4. J. Trilhe and G. Saucier, "WSI — The Challenge of the Future," *Proc. CompEuro 87: VLSI and Computers*, IEEE CS Press, Los Alamitos, Calif., Order No. 773, 1987, pp. 531-542.

5. R.O. Carlson and C.A. Neugebauer, "Future Trends in Wafer-Scale Integration," *Proc. IEEE*, Vol. 74, Dec. 1986, pp. 1,741-1,752.

6. J.F. McDonald et al., "The Trials of Wafer-Scale Integration," *IEEE Spectrum*, Vol. 21, No. 10, Oct. 1984, pp. 32-40.

7. W.M. Siu, "Monolithic VLSI vs. MCM: A Perspective on Performance, Yield, and Manufacturing," *Proc. IEEE Workshop Defect and Fault Tolerance in VLSI Systems*, IEEE CS Press, Los Alamitos, Calif., Order No. 2457, 1991, pp. 116-123.

8. K.A. Sack, R.C. Lyman, and G.T. Chang, "Evolution of the Concept of a Computer on a Slice," *Proc. IEEE*, Vol. 52, 1964, pp. 1,713-1,770.

GORDON·BELL·PRIZE
1992

1992 Gordon Bell Prize

For Outstanding Achievement in the Application of Parallel Processing to Scientific and Engineering Problems

Entries are due May 1, 1992, with finalists to be announced by June 30 and winners announced at the Supercomputing 92 conference in November 1992. Prizes of $1,000 each will be awarded in two of three categories: performance, price/performance, and compiler parallelization.

For more information,
contact Claire Azada,
IEEE Computer Society, (714) 821-8380.

Further reading

Apart from this special issue of *Computer*, several books and conferences have been dedicated to WSI. Of particular relevance is the annual IEEE International Conference on Wafer-Scale Integration. Readers looking for a historical perspective or for more on the latest research and development in WSI may find the following list helpful.

Conference proceedings

Proc. IEEE Int'l Conf. Wafer-Scale Integration, IEEE CS Press, Los Alamitos, Calif., 1989, 1990, 1991, 1992.

Proc. IFIP WG 10.5 Workshop on Wafer-Scale Integration (I, II, III), Elsevier North-Holland, New York, 1986, 1988, 1990.

Books

Wafer-Scale Integration, Jesshope, C., and Moore, W., eds., Adam Hilger, Bristol, England, 1986.

Wafer-Scale Integration, Swartzlander, E.E. Jr., ed., Kluwer Academic Publishers, Hingham, Mass., 1989.

Tewksbury, S.K., *Wafer-Level Integrated Systems: Implementation Issues*, Kluwer Academic Publishers, Hingham, Mass., 1989.

W. Kent Fuchs is an associate professor in the Department of Electrical and Computer Engineering, the Department of Computer Science, and the Coordinated Science Laboratory at the University of Illinois. His research interests include parallel computing and VLSI system design with emphasis on reliable computation.

Fuchs is the guest editor of the May 1992 *IEEE Transactions on Computers* special issue on fault-tolerant computing. He was technical program chair for the 1989 IEEE International Workshop on Hardware Fault Tolerance in Multiprocessor Architectures. He is currently on the Editorial Board of the *IEEE Transactions on Computers*.

Fuchs received a BS from Duke University in 1977, an MS from the University of Illinois, Urbana, in 1982, and a PhD from the University of Illinois in 1985, all in electrical engineering. He is a senior member of the IEEE and a member of the IEEE Computer Society.

Earl E. Swartzlander Jr. holds the Schlumberger Centennial Chair in Engineering in the Department of Electrical and Computer Engineering at the University of Texas at Austin. He is conducting research on application-specific processor technology and design, including high-speed computer arithmetic, systolic signal-processing architecture, VLSI and WSI technology, and system prototyping.

Swartzlander is editor-in-chief of *IEEE Transactions on Computers*, hardware area editor for *ACM Computing Reviews*, and the founding editor-in-chief of the *Journal of VLSI Signal Processing*. He was a member of the IEEE Computer Society Board of Governors from 1987 to 1991.

Swartzlander received a BS from Purdue University, an MS from the University of Colorado, and a PhD from the University of Southern California, all in electrical engineering. He is a fellow of the IEEE and a member of Eta Kappa Nu, Sigma Tau, and Omicron Delta Kappa.

Readers may contact Kent Fuchs at the Coordinated Science Laboratory, University of Illinois, 1101 W. Springfield Ave., Urbana, IL 61801, e-mail fuchs@crhc.uiuc.edu. Earl Swartzlander may be contacted at the Department of Electrical and Computer Engineering, ENS Bldg., Rm. 517, University of Texas at Austin, Austin, TX 78712, e-mail e.swartzlander@compmail.com.

ENTERPRISE-WIDE CASE FOR UNIX

THE **IDE**AL SOLUTION FOR C DEVELOPERS.

The C Development Environment is the first integrated CASE environment designed to work as well on your existing systems as it works on brand new projects.

It's the first CASE environment to support both incremental code generation *and* incremental design generation. You can take existing code, reverse engineer it, and place the information into a shared repository. Hit one button and generate publication quality documentation. Use point-and-click queries to access valuable information about code and design.

Developers can work on either design or code, depending on their preferences. And bring everything back in together without fear of overwriting one another. You're "synchronized." And you can "trust" the design and documentation to reflect the current state of the code.

The C Development Environment is an open solution that combines best-of-class tools from multiple vendors such as CenterLine, Frame Technology, or Interleaf. And, it's all backed by IDE training, consulting and support.

Why wait? Gain the benefits of this exciting technology on your current systems. The more you know, the more you'll want the award winning C Development Environment. Call us for more information and for a seminar schedule. 1 (800) 888-IDE1

A major technology and productivity breakthrough. Incremental design and code generation. For new and existing systems. Shared repository. Fast one-button documentation. Point-and-click queries. Immediate navigation among design, code, and documentation. An open solution integrating tools from multiple vendors. C Development Environment. There's nothing like it.

INTERACTIVE DEVELOPMENT ENVIRONMENTS

595 Market Street, San Francisco, CA 94105
Phone: 1(415) 543-0900 FAX: 1(415) 543-0145
E Mail: cdeinfo@ide.com.

© 1992 IDE. The IDE logo is a registered trademark and C Development Environment is a trademark of Interactive Development Environments, Inc. UNIX is a registered trademark of UNIX System Laboratories, Inc.

Reader Service Number 3

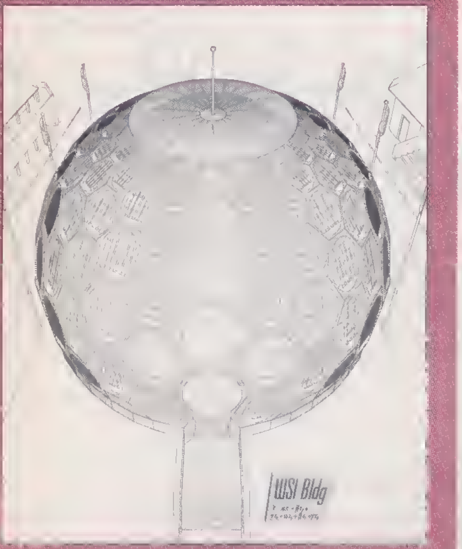

Task-Flow Architecture for WSI Parallel Processing

Robert W. Horst, Tandem Computers Inc.

This discussion explains the basics of task-flow architecture, examines a wafer-scale-integration implementation, and presents simulation results of a hypothetical machine.

The rapid evolution of integrated circuit technologies presents new challenges to computer designers. The increasing density of very large scale integration (VLSI) circuits has led to an exponential increase in the number of transistors available on each chip. Using full wafers for system fabrication can also increase the density per silicon device by two more orders of magnitude.

In planning ways to use the increased silicon area, designers tend to retain a familiar architecture while integrating more of the system on each device. Although this technique effectively reduces system cost, system performance increases at a modest rate. Dramatic increases in performance come only by exploiting the added parallelism possible with higher levels of integration. Although many current approaches to parallel processing define the problem in terms of selecting an optimal interconnection network for multiple microprocessors, this discussion proposes a new *task-flow* architecture that incorporates changes in the instruction set and programming paradigm to facilitate high-performance parallel processing using relatively simple interconnection networks.

This article discusses the basics of task-flow architecture, examines a possible implementation using wafer-scale integration (WSI), and presents simulation results of a hypothetical task-flow machine.

Task-flow architecture

The major difference between task-flow and other machines is the concept of sending computations to stationary data objects rather than sending data from memory to stationary processors. Task-flow machines do not contain separate processors and memories; instead, multiple cells contain both computing and memory elements. Data arrays are partitioned across multiple interconnected cells. Computation is performed by a set of tasks flowing through the network. Each task is executed by sending *transmission packets* (TPs) along linked lists of *memory packets* (MPs). Each MP contains a data element, the next instruction to

0018-9162/92/0400-0010$03.00 © 1992 IEEE

perform, and a link to the next MP. Each cell's memory is organized as sequentially addressable MPs. The linked-list nature of the MPs is useful in sparse matrix computations, database query processing, distributed reduction operations, associative memory, and coarse-grained application synchronization.[1,2]

Task flow is similar to the notion of *function request shipping* but is taken to the individual instruction level. The resulting machine has some similarities to hybrid dataflow architectures such as Monsoon and P-RISC, and to such other fine-grained machines as the J-Machine.[3] The ideas for flowing packets through the network have been partially derived from static dataflow machines.[4] The linear-array communications network was influenced by the WARP systolic array,[5] and the synchronization mechanism was adapted from the HEP processor.[6]

Although static dataflow machines may at first seem similar to task-flow machines, some fundamental differences exist. Pure dataflow machines take a functional view of processing in which no operations have "side effects" that modify permanent data structures. Task flow, on the other hand, employs a more declarative view of processing and associates a permanent modifiable storage location with every MP.

Traditional parallel architectures rely on a memory hierarchy in which frequently accessed data is migrated toward the processors. As illustrated in Figure 1, the traditional model implicitly assumes a round-trip delay to read memory data. The processor sends an address to the memory system, and data is returned to the processor. If the requested data values are not in the local cache, the round-trip delay can be substantial. Most machines block execution during cache fills while waiting for the remote data, although some systems have attempted to dispatch other tasks under these conditions.

To allow computations to flow efficiently between memory elements, a task-flow computer must support fine-grained computations. While some designers define fine-grained computation as 10 to 20 instructions, this project tried to provide efficient support for the execution of a single instruction at a processing node (cell). Unity grain size cannot be supported if a machine has separate instructions for communications and computation. Instead, a single instruction must be capable of performing an operation while continuing task execution at a remote node.

Figure 2 shows the most general task-flow system architecture. A network interconnects cells containing processing logic and memory. Multiple tasks can occur simultaneously, but unlike other parallel architectures, there is no one-to-one binding between tasks and processing elements. Tasks can flow through the network while executing a single instruction at each node or can remain local to one node throughout long computations.

Communicating sequential tasks perform the computations. Each task is sequential in time but is not necessarily sequential in space. Within a task, successive instructions can be located anywhere in the same cell or in a different cell. If most links point to MPs in the same cell, the instruction sequence is similar to a coarse-grained multiple-instruction, multiple-data (MIMD) machine. If all links point to an MP in the next cell, the instruction sequence resembles a linear systolic array. Varying MP placement can produce other communication patterns.

See the "Memory and transmission packets" sidebar for information on field organization.

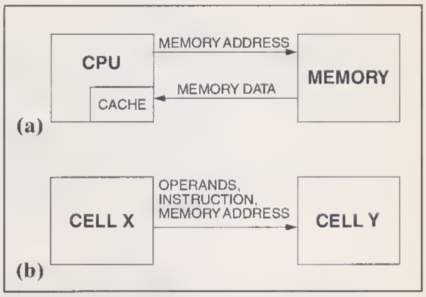

Figure 1. The round-trip memory-access delays required by conventional architectures (a) are avoided by the task-flow approach (b).

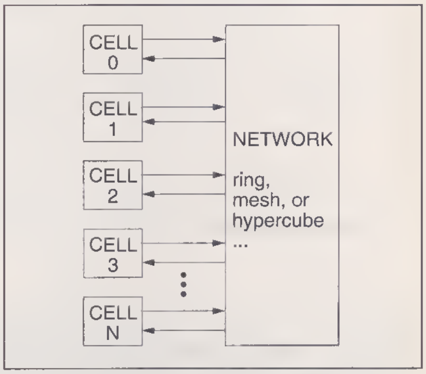

Figure 2. The most general organization for a task-flow machine has multiple cells and an interconnection network. Each cell contains computing and memory elements.

Memory and transmission packets

Within each cell, the memory is organized in *memory packets* (MPs), each comprising four fields:

Instr — The operation to perform at the next memory packet. This field also specifies instruction priority and whether to synchronize access to the next MP.

Link — A field linking the current MP to the next MP. The Link field contains the cell number and address offset within the cell of the next MP. Links are used for branch addresses and address pointers. Links can be dynamically read and written by variations of load and store instructions.

Data — Memory data. Storage for variables and instruction-immediate constants.

Lock — A 1-bit field used for synchronization. Similar to the HEP full/empty bit.[6]

Program execution is accomplished through the flow of *transmission packets* (TPs).These packets have the following fields:

Instr — The operation to perform at the destination MP.

Ra, Rb — Operand "registers" used to communicate intermediate results. Although these fields move along with the task, the programming language can read and write values to Ra and Rb in a way similar to register access in a sequential architecture.

Cell — The destination cell for this transmission packet.

Adr — The offset address of the next MP within the destination cell.

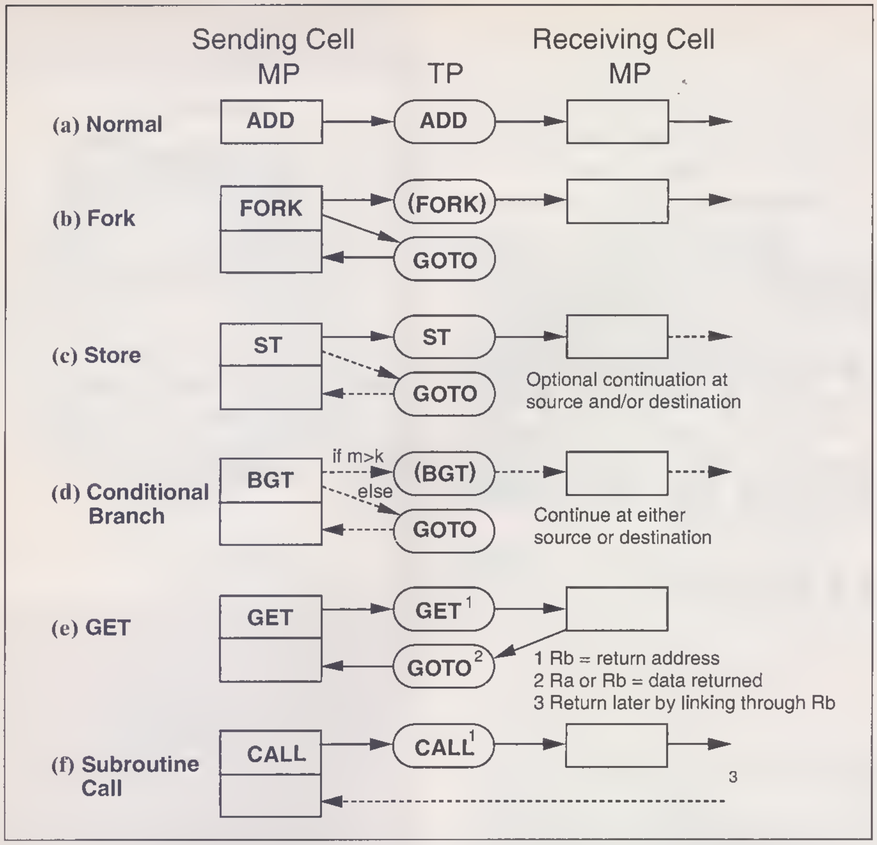

Figure 3. The six SWIFT instruction classes. The receiving MP can be in the same cell as (or in a different cell from) the sending MP. Some instructions complete execution at the sending cell; these instructions (in parentheses) are ignored at the receiving cell.

Task-flow simulation

SWIFT (Simulated Wafer-Scale Implementation of Flowing Tasks) is a register-transfer simulator that investigates the behavior of task-flow programs. SWIFT simulates a ring of cells with two pipeline stages between successive cells. In this simulation, transmission packets are sent in parallel in one clock period, although other implementations could multiplex TPs to conserve interconnection channels. (A more complete description of implementation and architectural alternatives can be found elsewhere.[2])

Each cell contains an arithmetic logic unit (ALU), a receive queue for holding incoming transmission packets, and a memory for storing MPs. Instructions operate on the two register fields of each TP and the addressed memory data. The TP at the head of the receive queue is decoded by a simulated programmable logic array (PLA) to begin execution. The PLA loads a busy counter that decrements each clock to simulate the time for fetching the addressed MP and executing the instruction. When the busy counter reaches zero, the instruction operation is performed and the MP instruction is decoded to determine how many packets should be generated: zero, one, or two.

Figure 3 illustrates the six SWIFT instruction classes. The instructions differ in the location and number of continuation tasks, and in the information carried in the register fields of the transmission packets.

In Figure 3a, the simplest instructions are executed by fetching an instruction and link from one cell, generating a transmission packet, and executing the instruction at the receiving cell. (Arithmetic instructions are executed in this manner.)

In Figure 3b, fork instructions generate two transmission packets. The sending cell dispatches one packet to the link address and a second to the next sequential MP in the same cell.

In Figure 3c, store instructions send one packet to store Ra or Rb in memory at the link address, and a second packet to continue execution at the next sequential address. Separate instructions store either the data field or link field of the addressed MP. While it is usually not desirable to continue the task performing the store, it is desirable to continue the sequential execution. The instruction set allows all four combinations of continuing or terminating the storing task and the sequential task.

Figure 3d shows how conditional branch instructions perform a test and then forward a TP to either the link address or the next sequential address. The conditional test is performed at the sending cell instead of the receiving cell. Branch tests are performed against the received value of register fields. If an ALU operation is sent to an MP containing a conditional branch, the branch tests the old register values. The concept is similar to the delayed branch used in many reduced instruction-set computers (RISCs), but an additional instruction must be inserted before rather than after the branch. An optimizing compiler would attempt to move an instruction between the modification of a register and a conditional branch based on that register.

In Figure 3e, the Get instruction sequence allows data to be retrieved from a target MP while ignoring the instruction and link fields of this target. The TP containing the Get instruction sends a return address in its Rb field. The receiving cell fetches the data and returns it to the Rb address. Get emulates the load instruction of a traditional architecture in that the address is sent to the data, the data returns, and execution continues at the next location. Without the Get instruction, a separate copy of each MP would be required for each unique algorithm accessing that data, because each MP contains only one instruction and one link. Algorithms requiring infrequent access to data can use Get, while those requiring the most frequent access can use the associated instruction and link fields to avoid the round-trip path to send the address and return the data.

In Figure 3f, the Call instruction resembles the first part of a Get. The sending instruction places the return address in Rb. At the destination, Call performs no operations other than fetching the next packet. At the end of the subroutine, a return operation is performed by specifying that Rb should be substituted for the link field. In hardware, the most significant bit of the MP link field indicates that Rb is to be used as the link.

Subroutines are reusable but generally not reentrant because local variables are held in the subroutine's MP data field. When a subroutine must respond to multiple callers, several copies of the subroutine can be created or the accesses to a single copy can be synchronized. The MP lock bits can be used to guarantee that multiple tasks executing the same subroutine do not interfere with each other when accessing local variables.

The instruction set is specified in 14 bits, including two 2-bit register-select fields, two 1-bit fields for destination and locking, a 6-bit opcode, and a 2-bit priority field. The priority field determines which TPs in the receive queue should be executed next. When TPs are received, those with higher priority are allowed to bypass any lower priority TPs already in the receive queue. Within a priority level, the oldest TP is executed first.

The simulator supports two methods for retrying instructions that attempt locked access to an MP variable. The simplest scheme is round-robin retry; instructions that fail return to the end of the receive queue. The second scheme, queue waiting, places failed TPs in the waiting state at the lowest priority level. When any instruction is successfully executed at the address that caused that instruction to be retried, the instruction is brought back to its original priority level and "promoted" back through the queue past any other waiting instructions of lower priority. For instance, if a locked LD (Load) instruction is sent to an MP before the value has been stored, the LD waits until the ST (Store) is successfully executed at that MP.

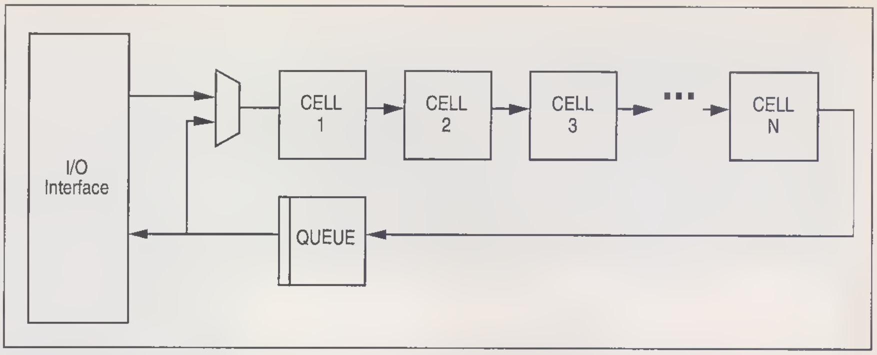

Figure 4. The unidirectional ring for intercell communication used by the SWIFT simulation. Such networks allow fast clock rates and wide data paths to provide high-bandwidth communications between neighboring cells.

Network design

Design decisions for SWIFT have been influenced by the desirability of implementing a WSI system. The usual difficulty in using an entire undiced wafer is that it contains a number of defects that cannot be allowed to disable the entire wafer. Many different wafer architectures have been proposed for configuring two-dimensional meshes and linear arrays from the wafer's functional cells. The difficulty with 2D meshes is that the regular interconnection pattern cannot easily be preserved through defective cells; configuring around the defects requires increasing some connections by several cell diameters. Another problem is that the maximum clock frequency in a 2D scheme has to allow for the worst-case distribution of defects.

Linear-array techniques sacrifice interconnect generality for improved nearest neighbor communication bandwidth. Linear-array WSI techniques generally offer a better harvest of good cells and require less configuration hardware than more complex 2D topologies. The applications currently being considered for SWIFT can be programmed so that nearly all communication traffic flows between packets local to one cell — or to neighboring cells. For these reasons, SWIFT currently assumes the topology of Figure 4, in which cells are arranged in a ring.

If the ring size grows much beyond 100 cells, the latency for nonlocal traffic can become unacceptable for some applications. Large systems containing many wafers can use a ring network within the wafer and a mesh or crossbar network between wafers. Multiple-wafer systems have not yet been studied in detail.

Cell architecture

Figure 5 illustrates the architecture of each cell. Incoming TPs arrive at the input pipeline register. The cell compares its local cell number with the des-

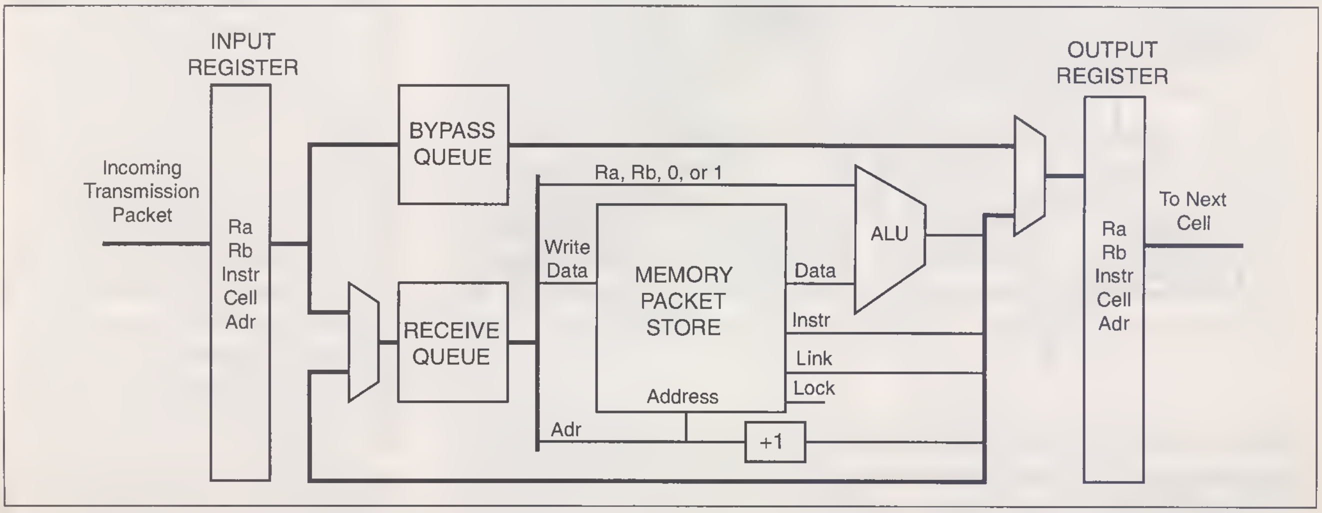

Figure 5. Simplified cell-block diagram.

Assembly language example

LIFT (Language Input for Flowing Tasks) is a high-level assembly language written to facilitate the development of task-flow programs on the SWIFT simulator. An example is illustrated by the flow diagram in Figure A and the code listing for a simple program shown in Table A. The program computes $A = B \ast C + B/D$ for $B = 6$, $C = 2$, and $D = 3$ using two communicating tasks. The rectangles contain MP labels used for both variable names and branch targets. The keyword MEM specifies an operand to be found in memory at the destination. Arcs show the instructions flowing between MPs. The tasks converge at BC, where one task stores the product and the other task adds the quotient to the stored product. The FILL and EMPTY keywords indicate that the lock bit at BC is to be honored to ensure that the store is performed before the ADD, regardless of which task reaches BC first.

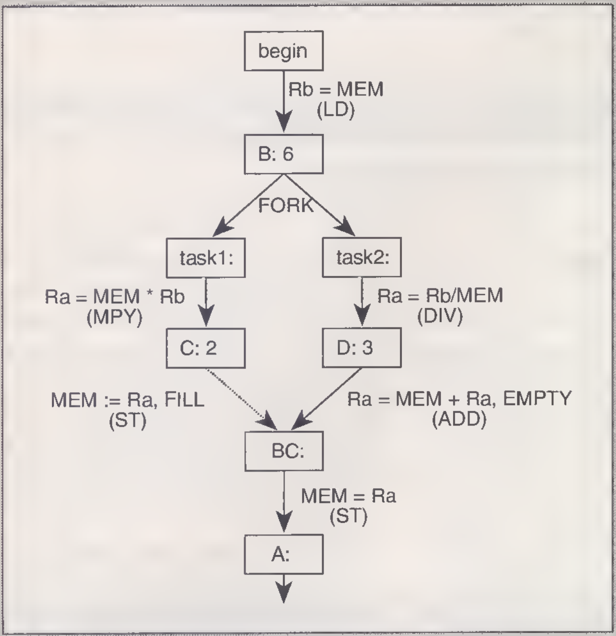

Figure A. Flow diagram for simple task-flow example.

Table A. Code listing for simple example program.

```
Current  Link to   Instruction to          Machine
MP Adr   next MP   perform at next MP      language

begin:   &B        Rb = MEM;               /* LD       */
B:=6     &task2    FORK;                   /* FORK     */

task1:   &C        Ra = MEM * Rb;          /* MPY      */
C:=2     &BC       MEM := Ra, FILL;        /* ST       */
                                           /* (HALT )  */

task2:   &D        Ra = Rb/MEM;            /* DIV      */
D:=3     &BC       Ra = MEM + Ra,          /* ADD      */
                     EMPTY;
BC:      &A        MEM = Ra;               /* ST       */
A: ;
```

tination cell of the TP. If the comparison indicates that the TP is bound for another cell, the TP passes through the bypass queue to the output pipeline register. When a completing instruction creates a new packet to be sent to another cell, the new packet has priority for the output register, and any bypass packets are queued.

Completing instructions can generate a number of continuation TPs: zero, one, or two. The destination of the link field can transmit a SND (send) packet to the current cell or the next cell. In addition, a SEQ (sequential) packet can be generated by a fork, store, or branch instruction. A SEQ packet directs a no-operation instruction to the next sequential location in the current cell.

When the received TP cell number matches that of the current cell, the TP propagates through the receive queue to the execution logic. When the cell is busy, several packets can be queued until the current instruction completes. The address field accesses memory while the received instruction is decoded. The data field of the addressed MP can be modified by a store instruction or can be used as an operand of a load or

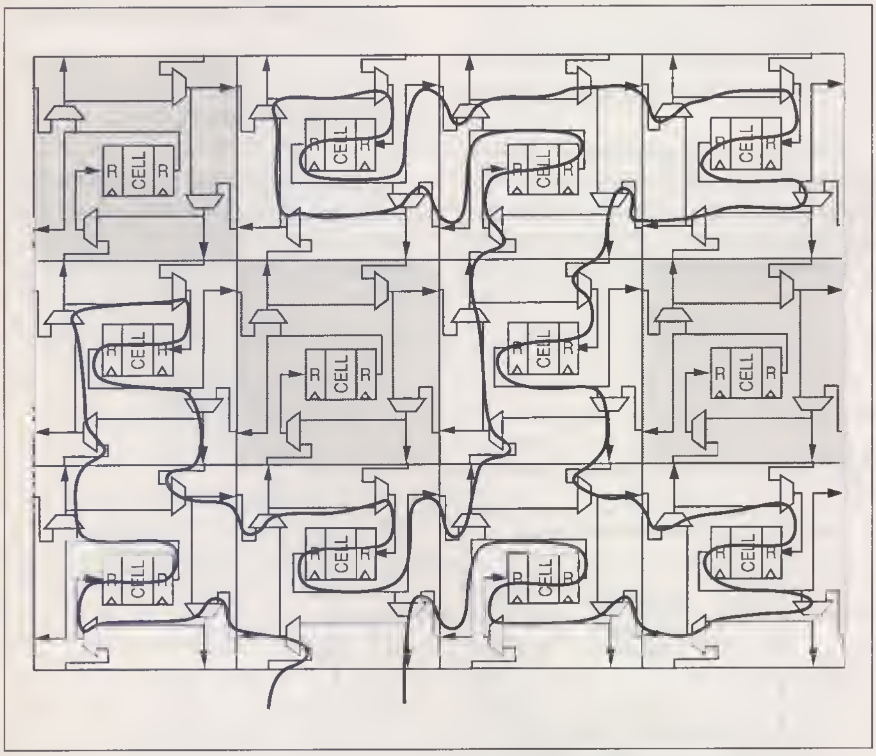

Figure 6. Wafer-scale architecture with each cell border connecting either to a neighboring cell or to an internal bus through a multiplexer. Defective cells can be avoided through properly setting multiplexer control latches (R = register, shaded = defective).

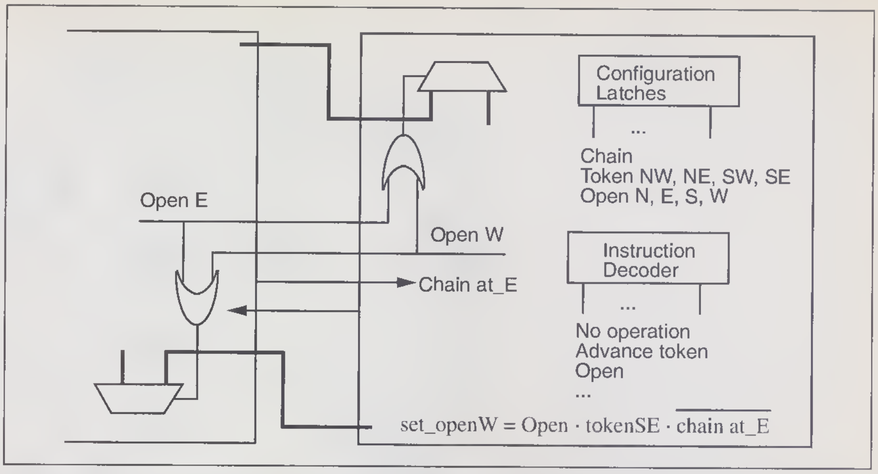

Figure 7. Wafer-configuration logic for in-circuit configuration (the dual-rail signals that prevent a faulty Open signal from corrupting the chain are not shown).

arithmetic instruction. Multiplexers select ALU inputs from the memory data, received registers, and constants. When the operation is complete, the output latch loads the outgoing TP. The decoded instruction in the output latch determines whether branch or fork instructions should generate SEQ and/or SND continuation packets. Comparing the current cell number and the destination cell number of the SND packet determines whether to continue the task locally or to send it to the next cell. If a retryable instruction is sent to an MP whose lock bit is set incorrectly, no new SND TPs are generated. Instead, the original TP recirculates to the receive queue to be retried later.

When resource conflicts are detected, flow control is provided by the bypass and receive queues at each stage. This process allows a few messages to be saved while a stall signal propagates backward through the pipeline. We have not yet investigated optimal queue sizes or deadlock avoidance schemes. Currently, the queue sizes are set to large values and statistics are recorded on the minimum queue size that gives full performance with no stalls.

See the accompanying sidebar for an example of the high-level assembly language developed to facilitate task-flow programs.

Wafer-scale architecture

WSI has been proposed for general-purpose processor interconnections as well as for such application-specific designs as neural network simulations.[7] WSI is particularly attractive for implementing the wide data paths required by task-flow machines. Keeping communications on-wafer reduces the latency of intercell transfers, and the high packaging density of WSI offers lower cost and more performance from a smaller volume.

Many potential task-flow applications seem to map well onto linear-array structures. Linear-array WSI architectures permit efficient harvesting of the working cells on a wafer, yet require less connection logic overhead than higher dimension networks. This reduction in configuration logic allows more cells per wafer for higher peak-processing performance.

Manning[8] and Catt and Aubusson[9] have proposed structures for linear arrays in which the possible connections are restricted so that each cell is part of one upstream and one downstream path. These schemes are less than ideal for processor connections because they do not support command pipelining. In addition, they require global signals for clocking and configuration. The proposed new WSI architecture solves these problems while reducing the amount of configuration logic and increasing the harvest of functional cells.

Chain WSI architecture

A new wafer-scale technique has been developed to allow linear arrays to be configured from the working cells on a partially good wafer.[10] The new *chain* architecture is particularly applicable to task-flow-machine implementations. The technique uses four multiplexers per cell to communicate with the four nearest neighbors. A simple algorithm sets multiplexer control latches to configure all working and reachable cells in a linear chain. Although Chang[11] has extended the chain architecture to three, six, or eight neighbors, this work assumes only the more typical case of four neighbors.

Figure 6 shows the configuration of a 12-cell section of a wafer. Identical cells are arranged in a checkerboard pattern with adjacent cells rotated by 180 degrees. The delay between pipeline registers in the ring always consists of four multiplexer delays and does not depend on defective cell distribution. The clock can be linearly distributed through the same multiplexers as the data to reduce clock skew and to tolerate defects in the clock distribution network. The WSI scheme supports fast clock rates due to the fixed delay between cells, the low-skew clock distribution scheme, and the lack of runtime arbitration and routing decisions that would be required of higher dimension networks.

Each cell has four 2:1 multiplexers, input and output pipeline registers (Rs), and a memory and arithmetic logic function. The multiplexers are controlled by configuration latches initialized during the power-on sequence. The multiplexers can be set to form a closed loop, or one or more multiplexers can be *opened* to form a larger loop with one or more neighboring cells. The cell is not limited to communicating with just two of its neighbors; setting the multiplexers appropriately allows communication with any 0-4 neighbors. Only 4 bits are required to configure the cell, and all 16 states are useful in different situations. Each functional cell serves as a link that can be added to a partially created chain. The chain is not required to have just one beginning and end; subchains can connect to any cell along the chain. The multiplexer configuration forms a single logical chain from the multiple subchains.

All register-to-register paths have exactly four multiplexer delays. This solves a problem inherent to many WSI designs — that is, the delay between pipeline stages is a constant rather than a function of defective-cell distribution. System performance is improved because the design does not have to allow for worst-case routing. Multiplexer and driver design can also be optimized because they always drive the same number of loads and amount of metal.

A diagram of the configuration logic is shown in Figure 7. Neighbors are

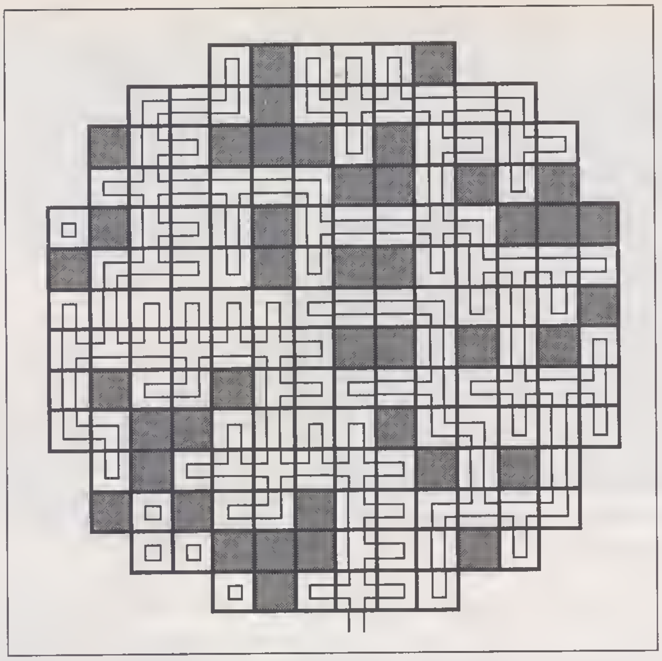

Figure 8. Configured 164-cell wafer with 75 percent cell yield.

indicated by abbreviations for the directions North, East, West, South, etc. At each border between two cells, multiplexers allow the cells to communicate. When either cell sets its Open signal for that border, both multiplexers are set to allow data to pass between the cells. (The Open signals are protected by a two-rail code to prevent a faulty Open from corrupting the chain.) Initially, a power-on reset signal sets all borders to be closed. Then, starting at the edge of the wafer, cells are opened until a chain is formed from all functional cells. Progress of the configuration algorithm is marked by a token passed along the chain. A chain bit is set to indicate that the cell has already been included as part of the chain.

Each wafer has several sets of bonding pads connected to a few cells around the chip periphery. At wafer-probe time, the tester identifies a set of pads connected to a nonisolated working cell, and these pads are bonded to the package. No further customization is necessary in the manufacturing process. The test and configuration algorithms are performed at wafer-test time or in the system during each power-up cycle. When power is first applied to the wafer, all cells are reset to be configured as loops with no external ports connected. The host system (or test system driving the wafer I/O pins) asserts the signal to open the port to the cell connected to the input pads. The token is then moved across the border and clockwise into a corner of the first cell. When the token first enters a cell, the cell opens a port to the neighbor clockwise from the token if that cell is not yet part of the chain. The host system then tests the newly connected cell and keeps the connection if the test is passed. If the test fails, the port is closed. Next, the token advances along the chain as it is currently constructed; the token moves clockwise around the cell until it finds an open port to another cell and moves into that cell. The algorithm repeats until the token has moved all the way around the chain, at which point all reachable cells are configured in the chain.

Figure 8 shows the routing of a 164-cell wafer in which 25 percent of the cells have been randomly assigned to be defective. In this example, 42 of the 164 cells are defective (shaded) and five are unreachable. The harvest is 117/122 = 96 percent.

Harvest and yield

A limited Monte Carlo simulation was run to predict yields for the 164-cell wafer. For simplicity, a Poisson distribution of defects was assumed, though it gives somewhat pessimistic results. Harvest is approximately 96 percent when the configuration logic yield is 75 percent. The cells not harvested are typically on the outer edge of the wafer because two of the four connection paths are unusable due to the wafer edge. Larger arrays would have even higher average harvests because the ratio of cells on the periphery to those in the interior would decrease.

Critical-area logic must work before a cell can be included in the chain. Simulations show that harvest drops sharply as the cell yield drops below 75 percent. The critical-area yield can be kept high by making the cells sufficiently small, by using partially functioning cells for routing, or by incorporating some defect tolerance into each cell. In a task-flow cell, the critical area includes clocks, pipeline registers, and control logic. The arithmetic logic need not be considered part of the critical area if the system is designed so that extra pipeline stages are allowed in the ring. The remaining cell area is devoted to packet queues and memory that can use spare rows and columns to improve yield.

With 0.6-micron technology, a full task-flow cell can be on the order of 100 square millimeters for a cell with approximately 1 megabit of memory. Assuming 2.5 defects per square centimeter, a cell with a critical area of less than 11.5 square millimeters is expected to yield 75 percent or more. Although we have not attempted any cell layouts, keeping the core logic to 10 percent of the cell area seems to be a reasonable goal.

Simulation results

The ultimate utility of task-flow computing rests on how well it performs on specific application programs. Simulation results of one application — sparse matrix-vector multiplication — illustrate some unique properties of task flow.

Sparse matrix operations are difficult to solve efficiently on parallel computers and supercomputers. Saad and Wijshoff report that performance is frequently as low as 5 to 20 percent of the performance achieved for dense computations.[12] The linked-list nature of task-flow machines is promising for improving the performance of sparse matrix computations. This section examines the performance of sparse matrix-vector multiplications on the SWIFT simulator.

Matrix-vector multiplication can be accomplished on a task-flow machine through a set of coordinated tasks. In one arrangement, each cell contains the elements from a set of matrix columns. Multiple tasks are started, with each task first multiplying a column by a vector element while leaving the products in memory. After completing the multiplications, each task begins to flow between the stored products to accumulate the sum for one row of the result vector. Multiplications are performed locally, while the additions are performed by tasks flowing through the network. Different data assignments

NEW CS PRESS TITLES

ARCHITECTURAL ALTERNATIVES FOR EXPLOITING PARALLELISM

edited by David J. Lilja

This book surveys the fine-grained parallel architectures that exploit the parallelism available at the instruction set level, the coarse-grained parallel architecture that exploits parallelism at the loop and subroutine levels, and the single-instruction, multiple-data (SIMD) massively parallel architectures that exploit parallelism across large data structures. The tutorial includes 37 articles discussing topics such as the potential of parallel processing for reducing the execution time of a single program, available parallelism in application programs, multiple-instruction issue architectures, and decoupled access/execute architectures.

Sections: Introduction, Fine-Grained Parallel Architectures, Coarse-Grained and Massively Parallel Architectures, Architectural Comparisons.

464 pages. February 1992. Hardbound. ISBN 0-8186-2642-9.
Catalog # 2642 — $80.00 Members $50.00

BRANCH STRATEGY TAXONOMY AND PERFORMANCE MODELS

by Harvey G. Cragon

This monograph presents a taxonomy that classifies and describes strategies, analytic models that permit the evaluation of each strategy under varying work load and pipeline parameters, and a modeling methodology that evaluates new branching strategies. It interprets analytic models, considers the implementation of parameters such as pipeline length and the location of branch-effective address ALUs, and investigates the various strategies and their performance models.

Sections: Introduction, Performance Modeling Method, Pipeline Freeze Strategies, Prediction Strategies, Instruction Sequence Alteration Strategies, Fetch Multiple Path Strategies, Composite Strategies.

120 pages. March 1992. Hardbound. ISBN 0-8186-9111-5.
*Catalog # 2111 — $40.00 Members $20.00***
(**Save $5.00 – special prepublication price)

GROUPWARE:

Software for Computer-Supported Cooperative Work

edited by David Marca and Geoffrey Bock

This book focuses on bringing computing face-to-face with groups of people as they work and discusses new software that enhances cooperation and augments human capability. *Groupware* concentrates on the task of designing software to fit the way groups interact in work situations and emphasizes the technical aspects in the development of software within the bounds of strong social and organizational factors.

Sections: Groups and Groupware, Conceptual Frameworks, Design Methods, Enabling Technologies – System-Related, Enabling Technologies – UI-Related, Computer Supported Meetings, Bridging Time and Space, Coordinators, What Makes for Effective Systems.

c.300 pages. April 1992. Hardbound. ISBN 0-8186-2637.
*Catalog # 2637 — $80.00 Members $40.00***
(**Save $10.00 – special prepublication price)

DIGITAL IMAGE PROCESSING, 2nd Edition

edited by Rama Chellappa

This collection is a revised and updated version of the first edition, published in 1985. The tutorial covers new research and advanced concepts and includes an overview of its five sections and a total of 47 papers of which 35 are new to this edition. It explores models for image data, filters in image processing, adaptive restoration of images, vector quantization, and image sequence compression. The text also contains a new section with six papers on emerging topics for digital image processing covering wavelets, morphology, neural networks, robust image processing, and high definition television (HDTV).

Sections: Image Models, Image Enhancement, Image Restoration, Image Data Compression, Emerging Topics, Bibliography.

816 pages. April 1992. Hardbound. ISBN 0-8186-2362-4.
*Catalog # 2362 — $75.00 Members $40.00***
(**Save $10.00 – special prepublication price)

ORDER TODAY! **Call toll–free 1–800–CS–BOOKS or *FAX* 714/821–4010**
(in CA call 714/821–8380)

***** Prepublication Prices Available To Members Only ! *****
You must return this card to get the prepublication discount.

ORDER FORM

Method of Payment:

☐ Check Enclosed

☐ VISA ☐ MasterCard ☐ American Express

Card # ______________________

Cardholder ______________ Exp. Date _____

Signature ______________________

IC #1

ORDER #	PRICE	QUANTITY	TOTAL

Name ______________________

Organization ______________________

Address ______________________

City ______________________

State ______________ Zip ______________

Phone No. ______________________

IEEE/IEEE Computer Society # ______________

Subtotal ______________

Please add applicable sales tax * ______________

Handling charge*** ______________

TOTAL $ ______________

Handling Charges***

Subtotal:	Add:
$ 1.01 to $ 50.00	$ 4.00
$ 50.01 to $ 75.00	$ 5.00
$ 75.01 to $100.00	$ 6.00
$100.01 to $200.00	$ 8.00
over $200.00	$15.00

Sales Tax*

CA residents 7.75%
CN residents add 7% GST
NY, NJ, and DC residents add applicable sales tax.

NEW CS PRESS TITLES

8th IEEE CONFERENCE on ARTIFICIAL INTELLIGENCE for APPLICATIONS (CAIA '92)

The proceedings contains 45 informative papers covering an application and/or an enabling technology that solves significant problems and stimulates the development of useful techniques that underlie real-world knowledge-based systems. It focuses on case studies and successful AI applications in business, industry, and government that rely on the knowledge of principles, techniques, and methodologies, as well as on use of the right tools. It also expands on state-of-the art developments in AI, exciting new approaches that are likely to affect future practices, and important new concepts for AI management.

Sections: Integrating Rule-Based Systems; Enabling Technologies; Diagnosis, Planning, Scheduling, and Constraint Satisfaction; Knowledge Acquisition and Learning; Intelligent Interfaces and Natural Language Processing; Case-Based Reasoning, Genetic Algorithm, and Fuzzy Logic; Design; Distributed Problem Solving and Real-Time Systems.

336 pages. February 1992. Softbound. ISBN 0-8186-2690-9.
Catalog # 2690 — $70.00 Members $35.00

8th INTERNATIONAL CONFERENCE on DATA ENGINEERING

Emphasizes the issues of technology transfer and focuses on the semantics and structuring of data in information systems design, development management, and use. The proceedings explores traditional applications as well as emerging technologies and issues in the area of automated data and knowledge management. Its 69 papers investigate AI and knowledge-based systems, application systems, tools and techniques, database management and structure, multimedia and distributed database systems, and integrity and security techniques.

Sections Include: Algorithms, Performance Techniques, Knowledge-Based Systems, Object-Oriented Databases, Database Tools, Data Integrity, Interoperability, Replicated Data Management, Performance Analysis, Storage Management, Logic and Databases, Object Management and Versioning, Deductive Databases, Query and Transaction Processing, Relational Database Techniques, Query Management, Database Modeling and Design.

680 pages. February 1992. Softbound. ISBN 0-8186-2545-7.
Catalog # 2545 — $100.00 Members $50.00

COMPCON Spring '92 —
37th Annual IEEE International Computer Conference

The 80 articles in this proceedings concentrates on the trends and technology that have led to the innovations of the 90s and the emerging technologies that will shape the industry in this decade. Its 12 tracks cover multimedia technology and applications, high-performance RISC systems, second generation RISC chips, recent microprocessors, high-performance networking technology and computing software, systems architecture, transaction processing, mobile computing, VLSI testing, and more.

Sections Include: Multimedia Standards and Authoring Tools, HDTV/Digital Video, Massively Parallel Computing, Distributed Memory Machines, SPARC Chips and Multiprocessors, Scalable PA-RISC Technology, High-Performance CMOS Microprocessors and Embedded Processors, Gbps Computer Interconnection and Networking, High-Speed Fiber Optic Products, Software Support for Parallel Processing, Advances in Testing.

520 pages. February 1992. Softbound. ISBN 0-8186-2655-0.
Catalog # 2655 — $90.00 Members $45.00

1992 INTERNATIONAL CONFERENCE on ROBOTICS and AUTOMATION (3–Volume Set)

This three-volume collection covers recent technical research in all aspects of robotics and manufacturing automation. The contributed papers in the proceedings delve into subjects such as:

- Modeling, performance evaluation, and simulation of dynamic systems
- Concurrent design of products and automated manufacturing
- High-level man/machine interfaces
- Robot sensing: vision, touch, range, force, sensor integration
- Multisensory perception and environment modeling
- Design and advanced control software for robotic mechanisms
- Task level planning, programming, reasoning, and skills acquistion
- Autonomous manipulation: reasoning, navigation, motion generation
- Multiple robot coordination and group robotics
- Teleoperation, telerobotics, and autonomous robots

c.3000 pages. April 1992. Softbound. ISBN 0-8186-2720-4.
*Catalog # 2720 — * $300.00 Members $150.00 (*prepublication price)*

To order : Call toll–free 1–800–CS–BOOKS or *FAX* 714/ 821–4010

PLACE
POSTAGE
HERE

IEEE Computer Society
Customer Service Center
PO Box 3014
Los Alamitos CA 90720-1264

change the mix of local and remote communications.

The simulated matrix has 1,000 rows and 1,000 columns, with 100,000 nonzero entries divided among the 100 cells. Each cell holds 1,000 entries with three MPs per entry. The simulation shows a speedup of 79 over the best sequential algorithm. The parallelism profile of Figure 9a shows that once the tasks are started, cell utilization remains near 100 percent except for the period when the tasks are finishing the multiplications and migrating to other cells for the first additions. The simulated times measure one pass of the matrix-vector product and do not include the time to initially load the matrix into the memory of the task-flow machine.

The communications profile in Figure 9b shows that most communications contact the next cell in the ring and that the traffic does not approach network capacity. The low communications requirements and large speedup show that the unidirectional ring network is appropriate for this problem. In one clock period, a maximum of 13 cells simultaneously send packets. The 2,500 clock periods do not use more than 5 percent of the total communications bandwidth, indicating that larger rings could be formed without bandwidth saturation.

Sparse matrix-vector multiplication is the primary operation required for the emulation of artificial neural networks. The 1,000 × 1,000 sparse matrix corresponds to a neural network with 1,000 neurons and 100,000 weights. The random connections are not restricted to any particular layered network structure, and simulation performance is not highly sensitive to the network topology as long as matrix elements are evenly partitioned among the cells. Each col-

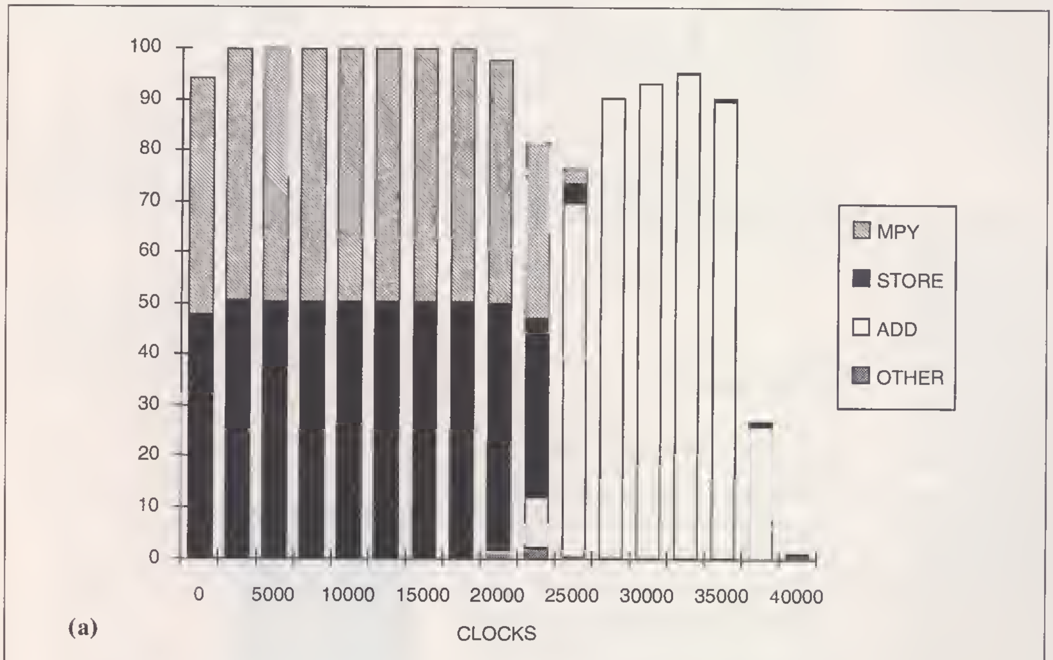

Figure 9a. Parallelism profile of the 1,000 × 1,000 sparse matrix-vector multiply on a 100-cell SWIFT. The vertical axis is a cumulative count of the number of cells executing each instruction type. The horizontal axis shows elapsed clocks through one pass of execution. The overall average speedup is 79 relative to sequential execution.

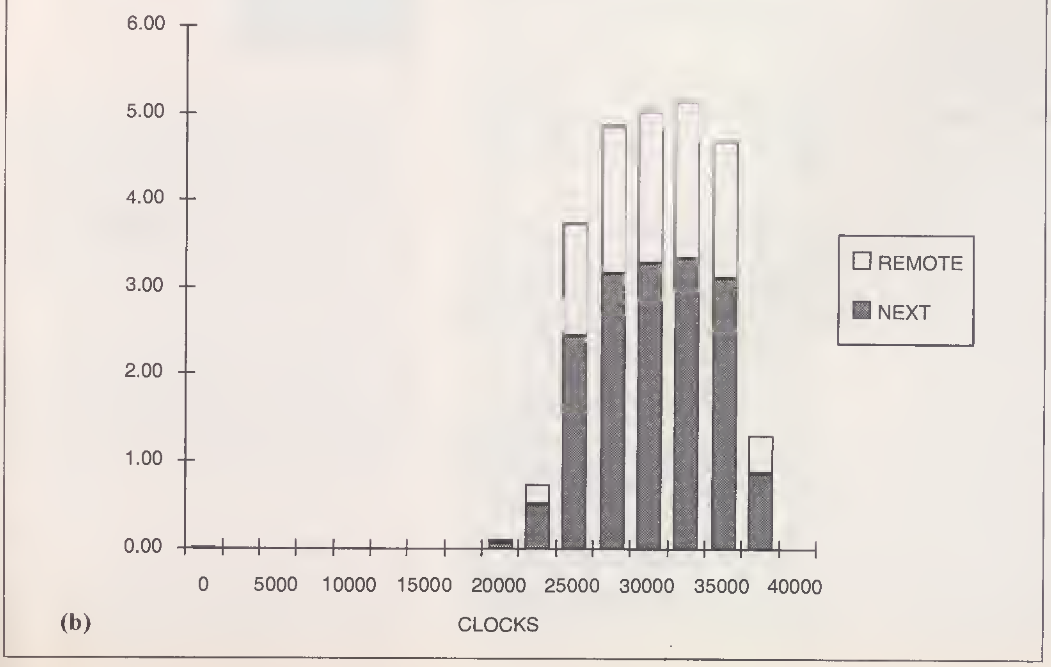

Figure 9b. Communications profile for the same simulation. The vertical axis shows the average number of cells sending packets to the next cell or another remote cell during each clock interval. During the period of maximum traffic, no more than 13 of the 100 cells simultaneously send an external packet.

umn of the weight matrix specifies the neuron fan-out, and each row specifies the fan-in. In this example, the random assignment of elements produced fan-in and fan-out with an average of 100 neurons, ranging from 68 to 137. The 100,000-connection network is a large network by artificial neural network standards but a very small one by biological standards.

Neural network performance (without learning) is generally measured in connections per second (CPS), where each connection is the multiply-accumulate associated with one element of the weight matrix. Assuming fixed-point arithmetic with five-clock integer operations and a 200-MHz clock, each pass of the SWIFT simulation requires 26,364 clocks, or 759 million CPS. Reducing precision and scaling to larger networks could improve these results.

Wafer-scale task-flow machines could address applications that are not well-suited to other parallel machines. Simulations have shown that flowing tasks can be effective in polling, task creation, distributed reduction operations, distributed synchronization, and matrix operations. Potential applications range from embedded controllers to signal processors to general-purpose multicomputers. The task-flow computation model is well-suited to classes of parallel computing that are not highly structured or vectorizable, such as sparse matrix operations, database searches, and neural network algorithms. The simple replicated cell structure gives high performance for a given silicon area and maps easily to WSI. The memory bandwidth scales with processing power. The system architecture allows the elimination of caches, giving some compensation for reduced memory densities.

Many variations of the basic architecture are possible. Memory space can be saved by eliminating unused MP fields, and performance can be increased by adding pipelining. Other investigative topics include extensions for virtual memory, high-performance I/O, and alternative interconnection networks.

A major remaining question concerns the programming environment for task flow. The declarative nature of task-flow programs does not require completely new programming languages, but program optimization requires either new programmer-visible constructs to allow manual assignment of MPs to cells or new compiler technology to determine good assignments automatically. Continued research will center on improving the software environment for the creation of task-flow programs. We hope to construct a prototype machine and evaluate performance on large-scale applications. ■

Acknowledgments

This work was performed at the Center for Reliable and High-Performance Computing at the University of Illinois with funding provided by Tandem Computers Inc. I thank Janak Patel and Andrew Chien for their insightful comments and suggestions on this work, and Bob Jardine and the anonymous reviewers for helping to improve this article.

References

1. R.W. Horst, "Task-Flow Computer Architecture," *Proc. Int'l Conf. Parallel Processing*, Vol. I, IEEE CS Press, Los Alamitos, Calif., Order No. 2101, 1990, pp. 533-540.

2. R.W. Horst, "Task Flow: A Novel Approach to Fine-Grain Wafer-Scale Parallel Computing," Coordinated Science Lab. Tech. Report CRHC-91-15, University of Illinois, 1991.

3. W.J. Dally, "The J-Machine: System Support for Actors," in *Towards Open Information Systems*, G. Agha and C. Hewitt, eds., MIT Press, Cambridge, Mass., 1990.

4. J.B. Dennis and D.P. Misunas, "A Preliminary Architecture for a Basic Data-Flow Processor," *Proc. Second Int'l Symp. Computer Architecture*, IEEE CS Press, Los Alamitos, Calif., Order No. M-096 (microfiche), 1975, pp. 126-132.

5. M. Annaratone et al., "The Warp Computer: Architecture, Implementation, and Performance," *IEEE Trans. Computers*, Vol. C-36, No. 12, Dec. 1987, pp. 1,523-1,538.

6. B.J. Smith, "Architecture and Applications of the HEP Multiprocessor Computer System," *SPIE*, Vol. 298, 1981, pp. 241-248.

7. F. Distante et al., "Fault-Tolerant Characteristics of the Linear-Array Architecture for WSI Implementation of Neural Nets," *Proc. Int'l Conf. WSI*, IEEE CS Press, Los Alamitos, Calif., Order No. 2126, 1991, pp. 113-119.

8. F.B. Manning, "An Approach to Highly Integrated, Computer-Maintained Cellular Arrays," *IEEE Trans. Computers*, Vol. C-26, June 1977, pp. 536-552.

9. R.C. Aubusson and I. Catt, "Wafer-Scale Integration — A Fault-Tolerant Procedure," *IEEE J. Solid-State Circuits*, Vol. SC-13, No. 3, June 1978, pp. 339-344.

10. R.W. Horst, "A Linear-Array WSI Architecture for Improved Yield and Performance," *Proc. Int'l Conf. WSI*, IEEE CS Press, Los Alamitos, Calif., Order No. 2013, 1990, pp. 85-91.

11. M.-F. Chang and W.K. Fuchs, "Design and Parallel Testing of Wafer-Scale Linear Arrays with High Harvest Rates," *Proc. Int'l Conf. WSI*, IEEE CS Press, Los Alamitos, Calif., Order No. 2126, 1991, pp. 285-291.

12. Y. Saad and H.A.G. Wijshoff, "A Benchmark Package for Sparse Matrix Computations," *Proc. Int'l Conf. Supercomputing*, ACM, New York, 1988, pp. 500-509.

Robert W. Horst is a technical director at Tandem Computers, where he has spent 11 years developing fault-tolerant computers. He is the principal architect of the NonStop Cyclone and NonStop TXP processors, and contributed to the Integrity S2 system design. His research interests include computer architecture, parallel processing, fault-tolerant computing, and wafer-scale integration.

Horst received a BS from Bradley University and an MS from the University of Illinois, both in electrical engineering. He recently completed a PhD in computer science at the University of Illinois. He is a member of the IEEE Computer Society and the ACM.

Readers can contact the author at Tandem Computers, 19333 Vallco Pkwy., Cupertino, CA 95014.

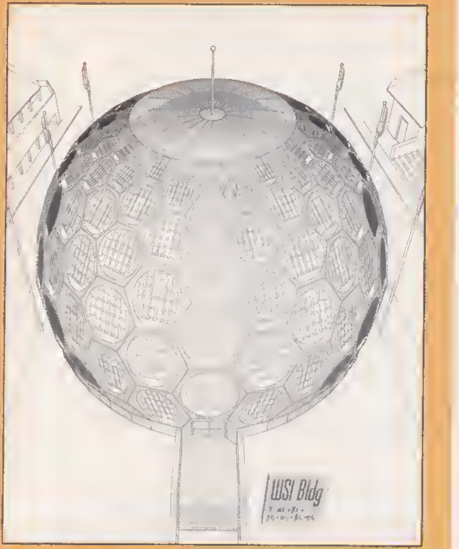

A Design and Yield Evaluation Technique for Wafer-Scale Memory

Koichi Yamashita, Fujitsu Laboratories

Shohei Ikehara, Fujitsu

Wafer-scale memory promises to provide greater yet cheaper storage volume than conventional memory systems using discrete chips. The authors used a Monte Carlo simulator for confirmation.

Wafer-scale memory, comprising full wafers and extra-large chips, has been reported in the literature[1-8] since it came under study in the 1970s. (See the "Wafer-scale memory research" sidebar.) Wafer-scale memory is expected to provide higher storage volume at a lower cost than conventional memory systems using discrete chips as wafer-scale memory technology is further developed.

In a wafer-scale integrated (WSI) system, it is important to optimize any redundant silicon area and maximize the number of functional circuits, because large redundant areas will decrease circuit density. The many papers that reported various yield analyses and simulation for very large scale integration (VLSI) circuits are also useful in a WSI design.[9-10] In addition, a WSI yield simulation has been designed for cellular designs.[11]

To evaluate the defect tolerance scheme for wafer-scale memory and predict the harvested capacity of the wafer memory, we developed a simulator that uses a Monte Carlo technique. This article describes the design of wafer-scale random-access memory that uses switching-register network logic.[12] In addition, we describe a simulator that selects the optimal defect tolerance scheme for our wafer-scale memory.

Architecture

The full-wafer memory system we designed consists of an array of chips, including memory core and additional network logic. We call our network logic a switching register (SW-REG). To isolate faulty areas and increase wafer memory capacity, each memory core block is partitioned, and the SW-REG and common signals are used to implement defect-tolerant schemes. We use transistor switches to reconfigure the wafer-scale memory; we don't use laser repair. The reconfiguration technique using transistor switches is the best method for mass production,

0018-9162/92/0400-0019$03.00 © 1992 IEEE

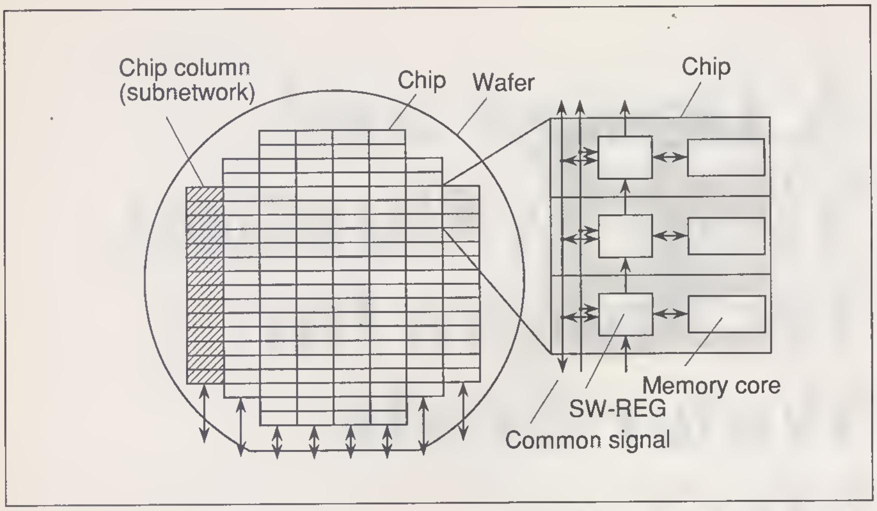

Figure 1. Wafer configuration.

since reconfiguring can only be done using conventional semiconductor processes and doesn't require special techniques or extra process steps. Furthermore, the same network logic can be used when the memory core size increases. As semiconductor process technology reduces pattern size, network logic will become smaller, decreasing the overall area overhead on the wafer. Area overhead is the ratio of SW-REG area to the whole chip.

Figure 1 shows the wafer configuration. Each chip is connected to its neighbors in the same column. The wafer network contains several subnetworks consisting of chip columns, each independent of the others. This architecture uses multiple bits and multiple data ports to improve data throughput and parallel access through interconnection lines to achieve faster access. However, this parallel-access architecture requires many common signals, so it is important to evaluate the defect tolerance schemes for these signals to obtain functional wafers.

Each column has 25 signals: the row address strobe (RAS), column address strobe (CAS), write-enable (WE), address (AD0-AD9), output enable (OE), data input/output (DI/O0-DI/O3), chip select address (CAD0-CAD4), scan clock (SCLK), and scan data (SDATA). For a wafer of eight columns, 200 signal lines are needed.

Each chip has an identification number and is selected for activation when the ID number is the same as that given by signals CAD0-CAD4. A memory address is given as AD0-AD9, translated by referencing external static RAM (SRAM), which contains information loaded from external electrically erasable programmable ROM (EEPROM) on defective memory blocks. The memory core has 4-bit data signals, as shown in Figure 2. The system accesses 20 bits of data from five chips in five different chip columns selected from eight columns when the wafer contains that number. The external multiplexer is required to select 20 out of 32 signals. When a chip of 100-nanoseconds access is used, the wafer system provides 600-ns access, including time for translating the address of a defective memory block. It takes about 500 ns to remap the address. When the system cycle is 100 ns, data throughput per data port is 40 Mbits per second, while that of the wafer system is 200 Mbits per second. Since only five chips are selected, the wafer uses little power. The power supply, control signals, and data signals can be placed on the same side of the wafer.

The SW-REG (see Figure 3) is additional logic for the wafer network, which consists of a switching register, a command decoder, and a decoder. The switching register, which comprises several registers connected serially, contains information about the chip ID and how to replace faulty circuits — as defined by SCLK and SDATA. It is programmed in an initialization period using SCLK and SDATA signals. No extra hardware is required in a wafer to set the ID number. Both chip address and command are sent serially to the switching register. The command decoder receives a command through the switching register when the signal CAD0-4 is the same as the chip ID (such as refresh for all chips in the same column). Several other inactivated chips are refreshed using the value given by signals CAD0-4 and the value held in the switching register. The decoder sends bypass control signals to the selector and changeover switches. The SW-REG receives signals CAD0-4 for the wafer network

Wafer-scale memory research

Nippon Telegraph and Telephone issued a report on a full-wafer 1-Mbit RAM and 4-Mbit ROM system fabricated by dynamically and automatically selecting good blocks out of duplicated blocks in a parity check during device power-on.[2,3] NTT research is now focusing on hierarchical redundant schemes.[4] Researchers in the European Strategic Programme for Research in Information Technology (ESPRIT) project have demonstrated the use of 4.5 Mbits on a 4-inch wafer using a 64-Kbit SRAM module.[5] The wafer is reconfigured with lasers and uses an additional nondiscretionary copper layer.

Two types of memory are commercially available. Inova Microelectronics has developed a 1-Mbit SRAM using 64-Kbit SRAM modules and an 8-Mbit SRAM using 256-Kbit SRAM modules.[6] An extra-large chip is fabricated using laser repair to bypass defects; it includes an additional nondiscretionary metal layer. Fujitsu and Anamartic fabricated the Catt's spiral proposed in the late 1970s[7] and developed a serial-access 200-Mbit wafer-scale memory using a 1-Mbit DRAM core.[8] The wafer system is linked by additional logic, forming a network configuration. It is designed to provide high wafer yield for a limited additional chip area by employing an interconnect scheme that isn't sensitive to the distribution of defects on a wafer. This memory system has an average 10-µs access time, and while faster than conventional magnetic systems, it is much slower than conventional memory chips.

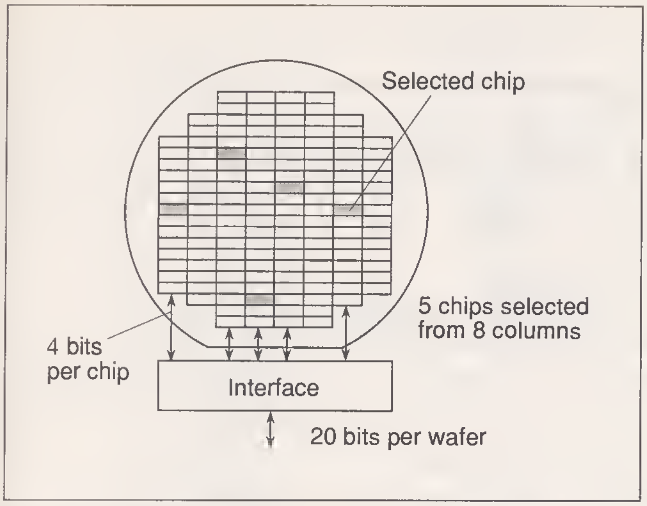

Figure 2. Wafer data access.

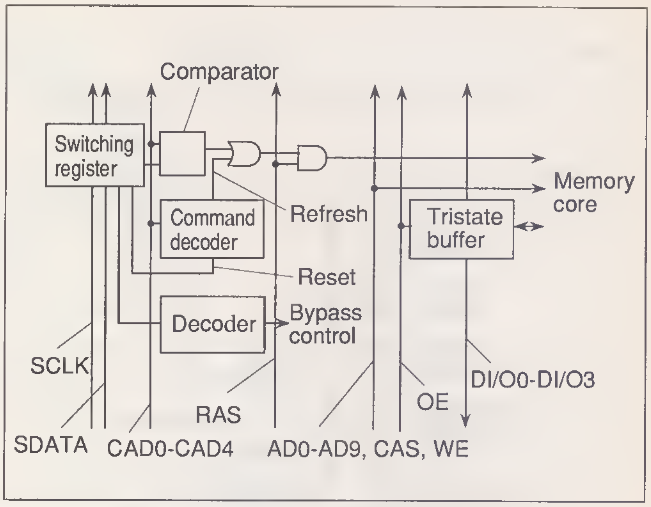

Figure 3. A switching register (SW-REG).

from the external peripheral circuit and accesses the memory core connected to it if it requires data access to the memory core. The SW-REG must also be bypassed to the next chip if the current one is defective. The SW-REG is estimated to require a little more than 1,000 gates, but will occupy only a small percentage of the entire chip area for a 4-Mbit dynamic RAM (DRAM) core — even if it contains the defect tolerance scheme described next.

Defect tolerance scheme

We studied several defect tolerance schemes that are combinations of drivers and spare lines for common signals, and bypass lines for the SW-REG.

Figures 4a and 4b show the defect tolerance schemes for common lines. A defect in a common signal causes an open- or short-line failure. Figure 4a also shows drivers for the common signals. The driver partitions the signal line into several parts, independent of defects in other parts, and decreases propagation delay on signal lines over several centimeters. Drivers also require very little area overhead. If there is no driver in the common signals, signals with defective lines can't be used. If common signals have drivers, the signal lines located prior to the defective part can be used. For propagation delay in real design, putting the driver into common signals every few chips is adequate.

Parameter	No driver	Driver
	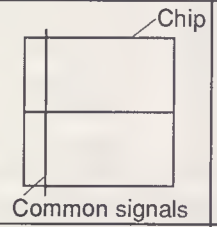	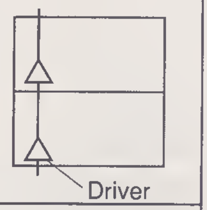
Additional area	Not required	Very small
Defect tolerance	None	Good

(a)

Parameter	No spare lines	Spare lines per chip column	Spare lines per chip
	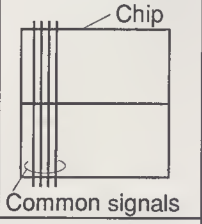	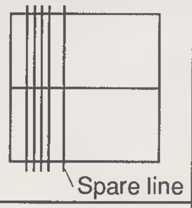	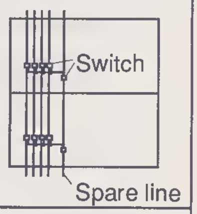
Additional area	Not required	Small	Large
Defect tolerance	None	Good	Very good

(b)

Figure 4. Drivers for common signals (a), and spare lines for common signals (b).

However, this evaluation assumes that each chip has a driver. Figure 4b shows spare lines for the common signals. When a chip column has spare lines for several common signals, a defective common signal line can be replaced by a spare line and the additional area is small. If a chip has one spare line for common signals, the column containing that chip is functional when there is a single faulty line but is defective if more occur. When a chip has spare lines and change-over switches for them, defective common signals are replaced by spare lines. Some defects in a chip column can be isolated, but the additional area is very large.

Parameter	No spare	Bypass line
	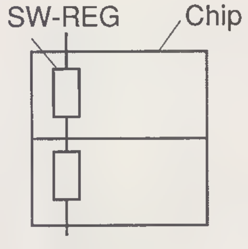	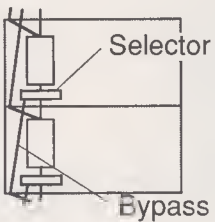
Additional area	Not required	Small
Defect tolerance	None	Good

Figure 5. Defect tolerance schemes for network logic.

Figure 5 shows defect tolerance schemes for the SW-REG. If a SW-REG is defective, the SW-REG without a bypass line can't be used; a defective SW-REG with a bypass line can be avoided at a small expense in the additional area. The next functional SW-REG selects a line to bypass the prior defective SW-REG. The SW-REG network using such a bypass remains functional if SW-REGs aren't consecutively defective. If two neighboring SW-REGs are faulty, the network becomes unusable. However, the probability of this occurring is very low. If the SW-REG is duplicated, an active SW-REG will be available, but the additional area is very large.

A combination of defect-tolerant hardware made up of the drivers and spare lines for common signals (see Figure 4) and a bypass line for a SW-REG (see Figure 5) constitute each scheme. We tested five different defect tolerance schemes, as shown in Figure 6. Common signals have drivers in all cases but case 1. The chip has no bypass line for the SW-REG in cases 1 and 2, but has one in cases 3, 4, and 5. There are no spare lines for common lines in cases 1-3, but there are spare lines in cases 4 and 5. Several spare lines are required for common signals. The spare lines are swapped for defective ones at the bonding pads. The area overhead is determined by designing a SW-REG layout that includes a defect-tolerant design. The area overhead required for cases 1-3 is 4.7 percent; for case 4, 5.1 percent; and for case 5, 13.4 percent. The SW-REG includes defect-tolerant circuits, test facilities, and a transistor switch for power supply.

Figure 6 also shows the usable chips in a column for each defect-tolerant condition if a defect occurs in a SW-REG or on the common signals. In this figure, the shadowed memory cores are unusable. In case 1, a short failure caused by a defect results in the entire column becoming nonfunctional, and an open failure results in all successive chips becoming nonfunctional. Case 1 has neither drivers nor spare lines for the common signals, and no bypass line for the SW-REG. In case 2, there are drivers but no spare lines for the common signals, and there is no bypass line for the SW-REG. When there is a defect in a common signal or in the SW-REG, all successive chips become nonfunctional. The defective area can be isolated by the driver. All SW-REGs prior to the defective ones are functional because the logic gates in the SW-REG can be isolated from the defective part. Case 3 has drivers but no spare lines for the common signals, and it has a bypass line for the SW-REG. A defect in the common signal makes all successive chips nonfunctional, while a defect in the SW-REG causes only the defective chip to fail. The last functional SW-REG sends a control signal to bypass the defective SW-REG to the selector. Cases 4 and 5 provide both drivers and spare lines for the common signals, as well as lines to bypass the SW-REG. There is one spare line per chip column in case 4 and one per chip in case 5. If a defect occurs in a common signal line, the line is replaced by an alternate. In case 4, a spare line is available for each chip column, and in case 5 a spare line exists for each chip (as shown in Figure 4). When these bypass lines are activated because of a defect in the SW-REG, they bypass the failed SW-REG and its connected memory core.

The power-supply network has a defect-tolerant scheme and can be separated from the power-fail SW-REGs by using the step-wire-bonding technique. The SW-REG has a transistor switch for the power supply to the memory core. If the transistor is turned off, the power-supply network won't supply power to the memory core.

Yield simulator

We developed a yield simulator to evaluate defect tolerance schemes and to calculate the usable circuit capacity. The simulator is implemented in C on a Unix workstation and uses X Window System, Version 11. We selected this computer environment for its portability. We entered the required information into the simulator database and started simulation by selecting the appropriate instruction. The database consists of chip, wafer, graphic, and parameter files. The input to the simulator involves the defect density, the defect distribution model (see sidebar, p. 24), and the defect tolerance scheme, plus the circuit area size for the wafer, chip, memory, and logic, and the failure type with probability weights for each part of the circuit. The memory capacity, number of memory partitions, and number of subareas are also input to the simulator. Each memory partition is the smallest usable unit of memory for a partially good chip. A memory partition containing a defective area can't be used, but it can be replaced by another partition having no defective area.

Figure 7 shows the simulator with several open windows and a pop-up

Case	Additional area (%)	Defect tolerance scheme			Defect in common signals	Defect in SW-REG
		Common signals		SW-REG		
		Driver	Spare lines	Bypass lines		
1	4.7	No	No	No	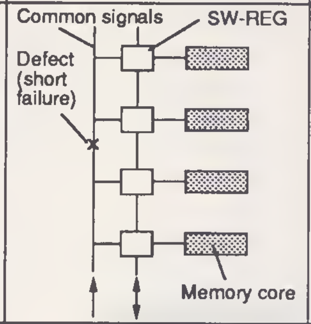 Common signals SW-REG Defect (short failure) Memory core	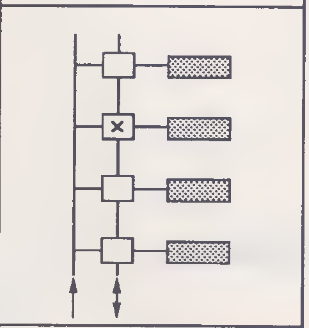
2	4.7	Yes	No	No	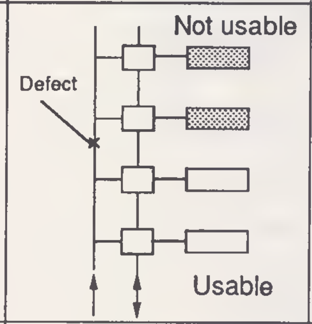	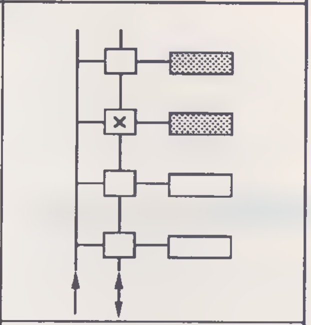
3	4.7	Yes	No	Yes	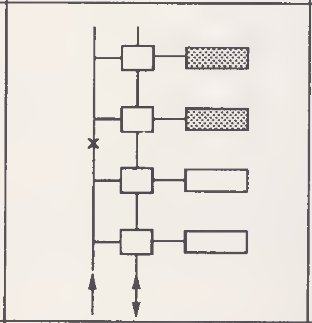	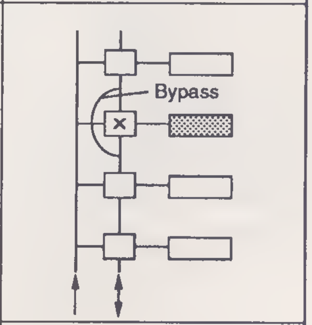
4	5.1	Yes	Spare line per column	Yes	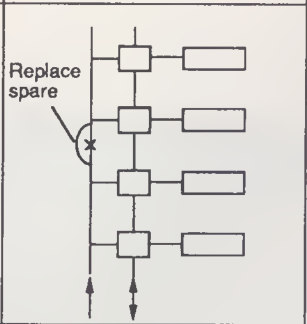	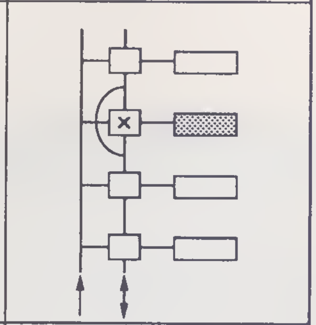
5	13.4	Yes	Spare line per chip	Yes	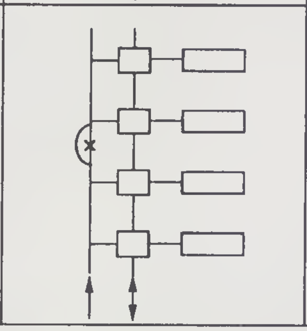	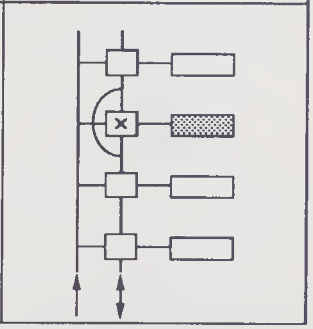

Figure 6. Variation of defect tolerance scheme.

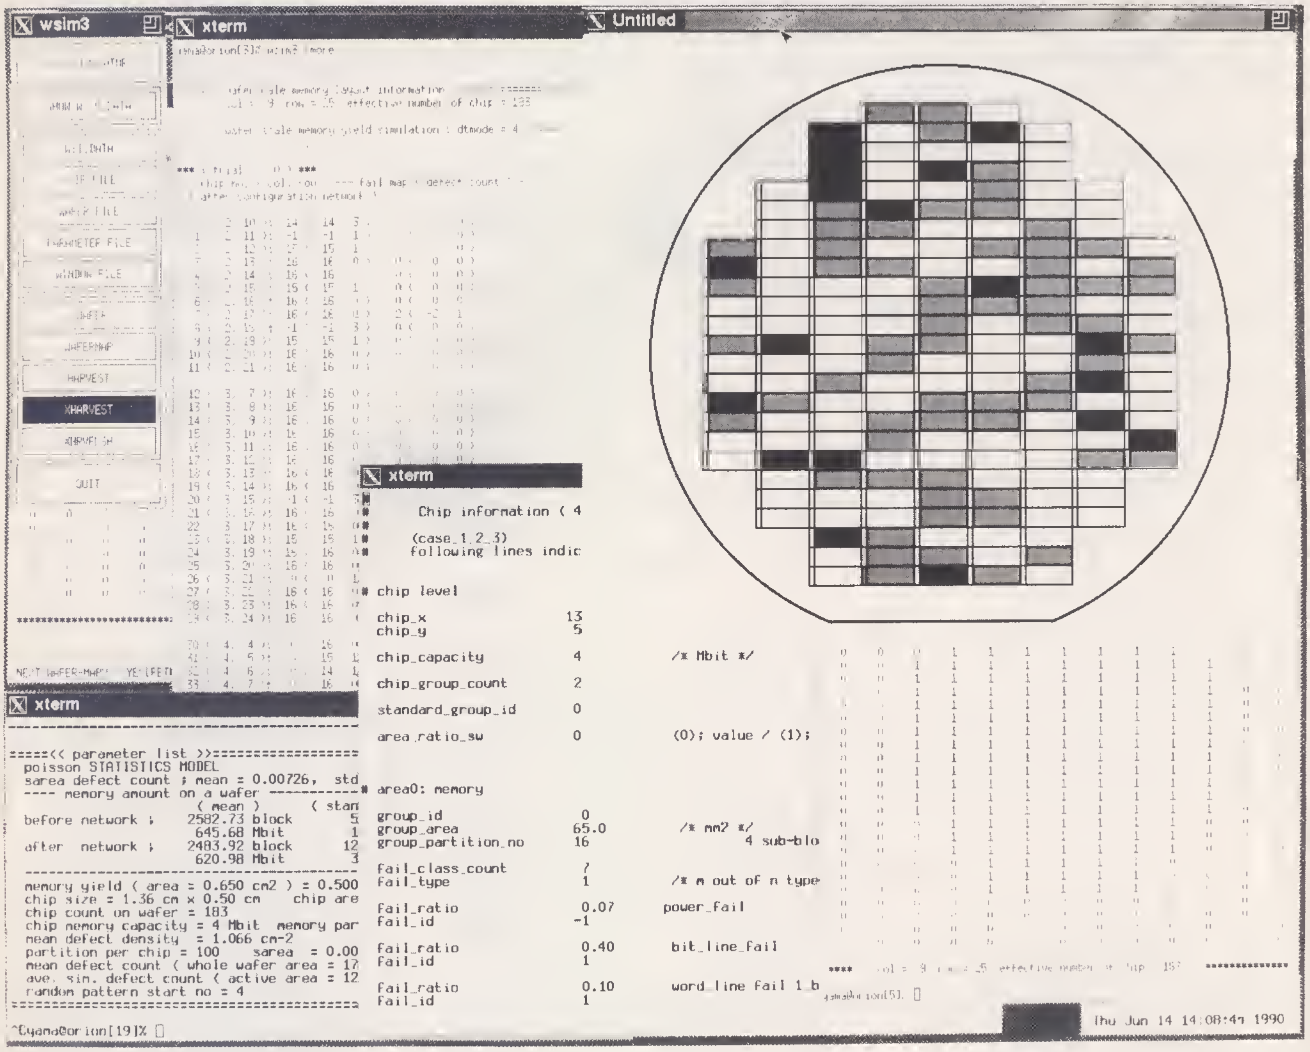

```
=====<< parameter list >>==================
  poisson STATISTICS MODEL
  sarea defect count ; mean = 0.00726,  std
  ---- memory amount on a wafer ------------
                         ( mean )       ( stan
  before network ;     2582.73 block         5
                        645.68 Mbit          1
  after  network ;     2483.92 block        12
                        620.98 Mbit          3
  -----------------------------------------
  memory yield ( area = 0.650 cm2 ) = 0.500
  chip size = 1.36 cm x 0.50 cm    chip are
  chip count on wafer = 183
  chip memory capacity = 4 Mbit  memory par
  mean defect density  = 1.066 cm-2
  partition per chip = 100    sarea  = 0.00
  mean defect count ( whole wafer area = 17
  ave. sim. defect count ( active area = 12
  random pattern start no = 4
==========================================

^Cyama@orion[19]%
```

```
      Chip information ( 4

     (case_1_2_3)
     following lines indic

chip level

chip_x                   13
chip_y                    5

chip_capacity             4          /* Mbit */

chip_group_count          2

standard_group_id         0

area_ratio_sw             0          (0); value / (1);


area0: memory

group_id                  0
group_area               65.0         /* mm2 */
group_partition_no       16                     4 sub-blo

fail_class_count          7
fail_type                 1          /* m out of n type

fail_ratio                0.07      power_fail
fail_id                  -1

fail_ratio                0.40       bit_line_fail
fail_id                   1

fail_ratio                0.10       word_line fail 1_b
fail_id                   1
```

Figure 7. The simulator screen.

Defect distribution model

The number of defects based on a given distribution model can be established as shown below.

When the Poisson model is used, the number of defects in each subarea is determined using Poisson random numbers obtained by executing steps 1 to 4, below. The cumulative distribution function $F(k)$ of the Poisson distribution is

$$F(k) = P\{N \le k\} = \sum_{j=1}^{k} \frac{e^{-\lambda} \cdot \lambda^{j}}{j!} \tag{1}$$

where λ is the mean defect number of the subarea.

Step 1: Calculate $F(k)$, where k is 0, 1, ... M. (M is the largest number of k that is a defect number of the subarea and is determined by defect density and subarea size.)

Step 2: Set k to 0 ($k = 0$).

Step 3: Generate a uniform random number R that is a numerical value between 0 and 1.

Step 4: Set N to k, N being a Poisson random number if $F(k)$ is larger than or equal to R.

Set k to $k + 1$ and go to step 4 if $F(k)$ is smaller than R.

When the generalized negative binomial (GNB) model is used, the cumulative distribution function $F(k)$ is

$$F(k) = P\{N \le k\} = \sum_{j=1}^{k} \frac{\Gamma(\alpha + j)}{j! \cdot \Gamma(\alpha)} \frac{\left(\frac{\lambda}{\alpha}\right)^{j}}{\left(1 + \frac{\lambda}{\alpha}\right)^{j+\alpha}} \tag{2}$$

where α is a clustering parameter and Γ represents the gamma function. The procedure in the GNB model parallels that of the Poisson model.

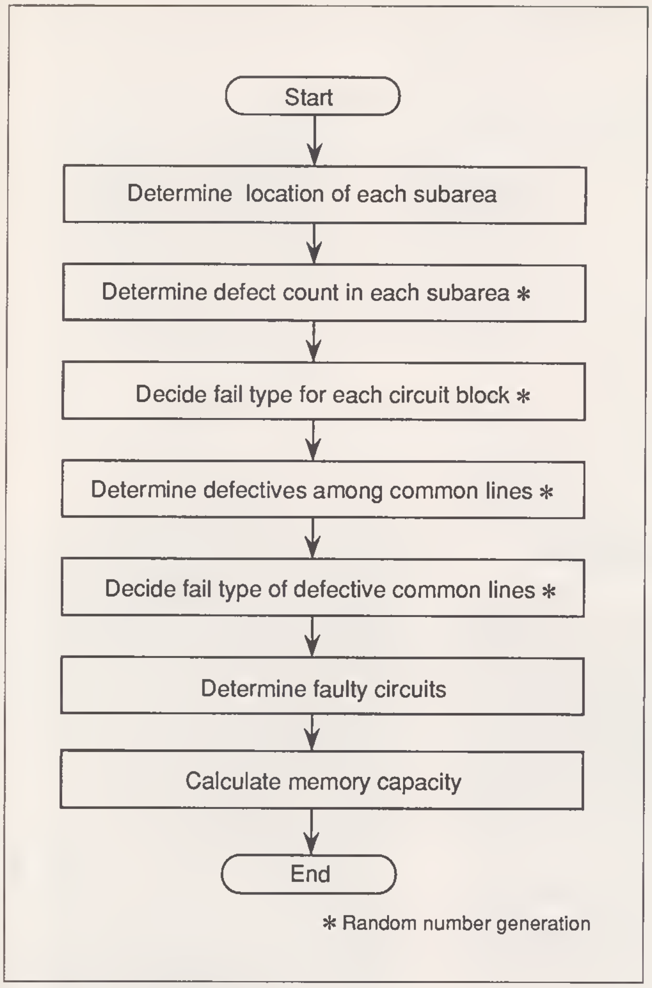
Start
Determine location of each subarea
Determine defect count in each subarea ∗
Decide fail type for each circuit block ∗
Determine defectives among common lines ∗
Decide fail type of defective common lines ∗
Determine faulty circuits
Calculate memory capacity
End
∗ Random number generation

Figure 8. Simulation flow.

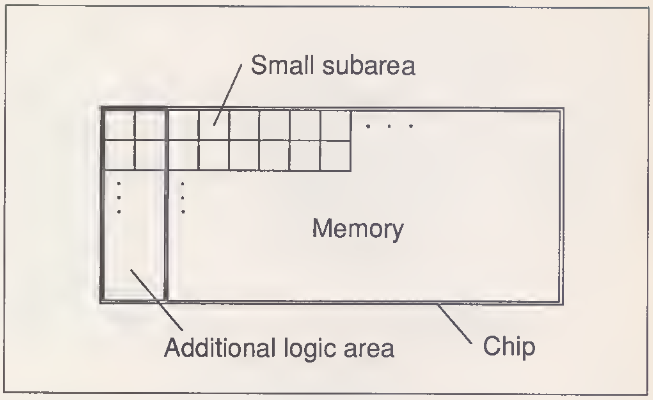

Figure 9. Location of minute area.

Subarea	Defects	Circuits
0	0	Additional logic area
1	0	
2	1	
3	0	
4	0	Memory
5	0	
6	2	
• • •	• • •	
99	1	

Figure 10. Determination of defects in circuits.

menu for invoking the simulation. The windows show the simulation file name, the simulation results (in detail or summary), and a map of the wafer (including failures) after configuration. The simulator can also calculate the wafer's effective memory capacity. For example, if XHARVEST is selected, the simulator calculates and shows the number of usable circuits, and displays a new window showing good, partially good, or defective chips on the wafer (at upper-right in Figure 7).

Figure 8 shows the harvest simulation flow. The simulation system sections the chip into very small subareas and assigns circuits to the appropriate locations. The simulator randomly generates defects using the defect density and distribution function and determines which defects cause faults on the chip. The simulator can use a cumulative Poisson distribution function or a generalized negative binomial (GNB) model.[9] The simulator configures the network on a wafer, counts the usable memory blocks, and calculates the overall memory capacity of the wafer. In Figure 8, four different random-number series are used in the program steps marked by asterisks.

As shown in Figure 9, a chip is first divided into small subareas covering the memory area and the networking logic area. The simulator then determines which circuit block each subarea is placed in by using the area ratio of the circuit block written in a database chip file. The chip is divided into a numerous subareas for simulation. The simulator determines the number of defects in a subarea, on the basis of a given defect distribution model.

Figure 10 presents an example of the defects generated for a chip divided into 100 subareas; the area ratio determines that four subareas contain the logic area, while 96 subareas contain memory. In this example, subareas 0 and 1 have no defects, subarea 2 has one defect, and so on. Defect generation results in one defect in the additional area and several others in the memory area. After determining the number of defects in the circuits, the simulator determines the type of failure caused.

The simulator then specifies the failure type (such as power, bit-line, or word-line). It also uses random numbers and failure probabilities read from a chip file. Figure 11a contains six types of failures — bit-line, word-line, single-bit, multibit, block, and chip — in the memory area. The database contains the occurrence probability of each type of failure as well as the number of defective memory partitions. The simulator selects the failure type by generating a random number between 0 and 1. For example, if the generated number falls in the bit-line region (the left-most re-

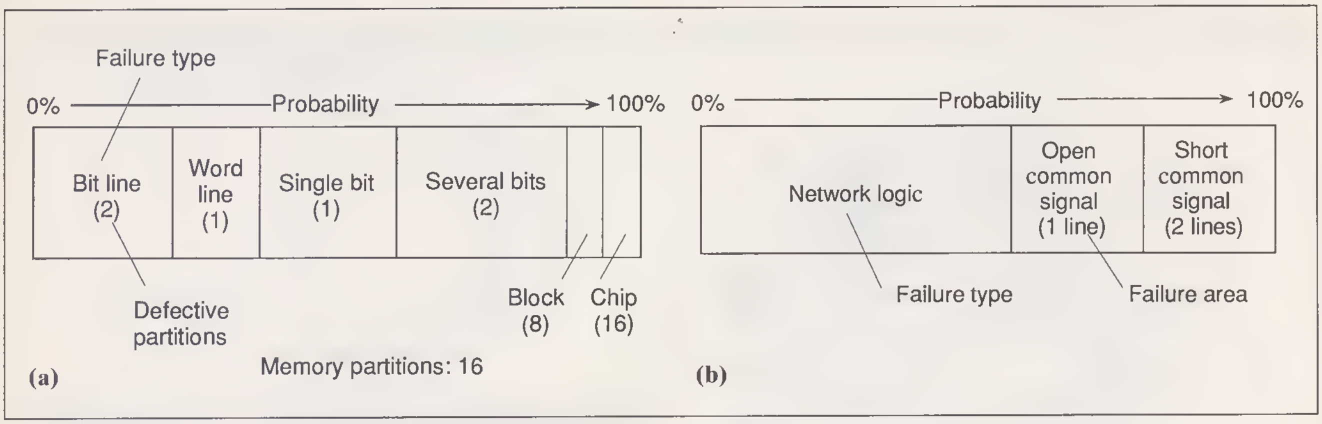

Figure 11. Memory area failures (a), and logic area failures (b).

gion in Figure 11a), then a bit-line failure occurs and two memory partitions fail. The numbers in parentheses in Figure 11a represent the defective memory partitions for memory partitioned into 16 blocks. The type of failure that occurs in the logic area is determined by a method similar to that for the memory area. Figure 11b shows three failure types — additional network logic, open common-signal, and short common-signal failures — in the logic area. When an open common-signal failure occurs, this means a single line failed. The short common-signal failure causes two defective lines.

Then, on the basis of the defect tolerance scheme in the parameter file, the simulator decides whether the wafer-scale system can replace a defective circuit with a spare. Finally, the simulator configures the network on a wafer, counts usable memory blocks, and calculates the overall memory capacity of the wafer.

Simulation results

Using our simulator, we optimized the fault tolerance scheme — including redundant circuits and lines in the networking logic and the number of memory partitions. This article shows the simulation results using only the Poisson distribution. The values of all input parameters were determined using real circuit sizes and experimental values on production lines for several months. The experimental values at higher chip yields were distributed on the Poisson model.

Figure 12 shows the capacity of a 6-inch wafer, comparing the five defect tolerance cases. The simulation uses a 4-Mbit DRAM core as the basic memory partitioned into 16 blocks. In cases 1-4, wafer memory increases with the increasing defect tolerance; that is, case 4 is better than case 3, case 3 is better than case 2, and case 2 is better than case 1. The capacity in case 3 is comparable to that of case 4 because there are few defective common signals when the memory core yield exceeds 50 percent. In case 5, the memory capacity begins to fall off at the higher core yields, mainly due to the large additional area needed to implement this case. Overall, case 4 is the best, due to the SW-REG bypass lines and alternate paths in each chip column for common signals. We get about 600 Mbits when the memory core yield is 54 percent.

Figure 13 compares memory capacity with the number of partitions for the defect tolerance scheme in case 4. When memory core is partitioned into two to eight blocks, wafer memory capacity increases with the number of partitions. When memory core yield is low, the capacity gradually increases with the number of partitions, even if the number exceeds eight. However, when a memory core yield is high, the capacity becomes saturated. Therefore, taking into account hardware overhead need-

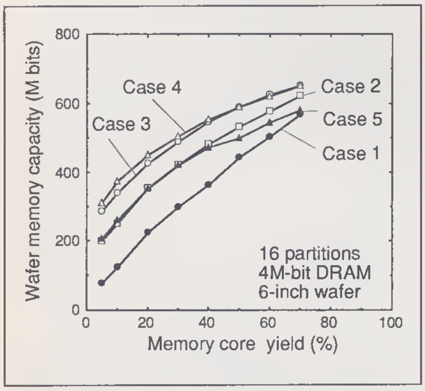

Figure 12. Comparison of the five defect tolerance cases.

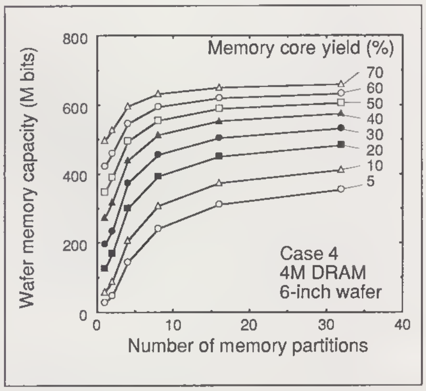

Figure 13. Comparison of memory capacity with the number of memory partitions.

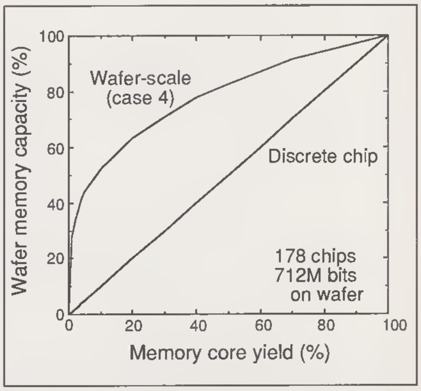

Figure 14. Comparison of wafer-scale memory capacity with discrete-memory capacity.

ed for translating the address of a defective memory block, the memory core should ideally be partitioned into eight to 16 blocks.

Figure 14 compares wafer-scale memory and discrete-memory capacities for different memory core yields when both chip counts are the same. Wafer-scale memory provides much more capacity than discrete chips because it can use the memory of partially functional chips. The full capacity of wafer-scale memory with a 5 percent memory core yield can be about eight times greater than the equivalent number of discrete chips in a wafer. For a 50 percent memory core yield, the wafer-scale memory capacity can be up to 1.6 times higher. There are fewer chips on a wafer for wafer-scale memory than for discrete memory because of the area overhead for the SW-REG. The curve for wafer-scale memory will cross and go below that for discrete chips at high memory core yields.

Table 1. Wafer-scale memory specifications.

Feature	Specification
Memory capacity	600 Mbits
Memory core	4-Mbit DRAM
Wafer size	6 inches
Data throughput	200 Mbits/sec.
Access time	600 ns
Data access	20-bit parallel

Wafer-scale memory using a 4-Mbit DRAM core is summarized in Table 1. Such memory has hundreds of lines so that it can achieve parallel access for high data throughput. Estimating the defect tolerance scheme gives us an area-efficient system with optimized additional circuits.

In our system, access is 20 times faster than in serial-access memory, and throughput is 10 times higher.[8] This parallel-access architecture for wafer-scale memory is thus advantageous for high-performance, high-density, and low-cost memory storage. ■

Acknowledgments

We thank Takahiko Misugi, Hajime Ishikawa, Gensuke Goto, Shinpei Hijiya, Fumio Baba, Takeo Tatematsu, and Masashi Nagashima for their valuable contributions and support.

References

1. J.F. McDonald et al., "The Trials of Wafer-Scale Integration," *IEEE Spectrum*, Vol. 21, No. 10, Oct. 1984, pp. 31-39.
2. Y. Egawa, N. Tsuda, and K. Masuda, "A 1-Mbit Full Wafer MOS RAM," *ISSCC 79, Digest Technical Papers IEEE Int'l Solid-State Circuits Conf.*, Lewis Winner, Coral Gables, Fla., 1979, pp. 18-19.
3. Y. Kitano et al., "A 4-Mbit Full Wafer ROM," *IEEE J. Solid-State Circuits*, Vol. SC-15, No. 4, Aug. 1980, pp. 686-693.
4. N. Tsuda, "A Defect- and Fault-Tolerant Design of WSI Static RAM Modules," *Proc. Int'l Conf. Wafer-Scale Integration*, IEEE Computer Soc. Press, Los Alamitos, Calif., Order No. 2013, 1990, pp. 213-219.
5. B. Nasreddine et al., "A Reconfigurable SRAM 4.5-Mbit WSI Memory," *Defect and Fault Tolerance in VLSI Systems*, Plenum Press, I. Koren, ed., Vol. 1, Oct. 1988, pp. 227-242.
6. R. Bourassa, T. Coffman, and J.E. Brewer, "Progress in WSI SRAM Development," *Proc. Int'l Conf. Wafer-Scale Integration*, IEEE Computer Soc. Press, Los Alamitos, Calif., Order No. 2013, 1990, pp. 13-19.
7. R. C. Aubusson and I. Catt, "Wafer-Scale Integration — A Fault-Tolerant Procedure," *IEEE J. Solid-State Circuits*, Vol. SC-13, No. 3, June 1978, pp. 339-344.
8. N. MacDonald et al., "200 Mbit Wafer Memory," *ISSCC 89, Digest Technical Papers IEEE Int'l Solid-State Circuits Conf.*, John H. Wuorinen, Castine, Maine, 1989, pp. 240-241.
9. C.H. Stapper, F.M. Armstrong, and K. Saji, "Integrated Circuit Yield Statistics," *Proc. IEEE*, Vol. 71, No. 4, Apr. 1983, pp. 453-470.
10. H. Walker and S.W. Director, "VLASIC: A Catastrophic Fault-Yield Simulator for Integrated Circuits," *IEEE Trans. Computer-Aided Design Integrated Circuits and Systems*, Vol. CAD-5, No. 4, Oct. 1986, pp. 541-556.
11. J.C. Harden, J.J. Wang, and B.W. Tebbs, "A High-Level WSI Yield Simulation System," *Proc. Int'l Conf. Wafer-Scale Integration*, IEEE Computer Soc. Press, Los Alamitos, Calif., Order No. 1901 (microfiche only), 1989, pp. 235-243.
12. K. Yamashita et al., "Evaluation of Defect Tolerance Scheme in a 600-Mbit Wafer-Scale Memory," *Proc. Int'l Conf. Wafer-Scale Integration*, IEEE Computer Soc. Press, Los Alamitos, Calif., Order No. 2126, 1991, pp. 12-18.

Koichi Yamashita joined Fujitsu Laboratories, Kawasaki, Japan, in 1982 where he performs VLSI design research. His research interests include wafer-scale integration, defect- and fault-tolerant design, yield enhancement techniques, and rapid prototyping architecture for application-specific integrated circuits.

He received BE and ME degrees in electrical engineering from Kobe University in 1980 and 1982, respectively. He is a member of the IEEE Computer Society and the Institute of Electronics, Information, and Communication Engineers of Japan.

Shohei Ikehara joined Fujitsu Limited, Kawasaki, Japan, in 1975 to help develop the company's mainframe design.

He received the BE degree in electrical engineering from the University of Electro-Communications in Tokyo in 1975. He is a member of the Information Processing Society of Japan.

Yamashita can be contacted at Fujitsu Laboratories Limited, ULSI Systems Laboratory, 10-1 Morinosato-Wakamiya, Atsugi 243-01, Japan, e-mail yamashita@flab.fujitsu.co.jp. Ikehara can be contacted at Fujitsu Limited, Mainframe Division, 1015 Kamikodanaka, Nakahara-ku, Kawasaki 211, Japan.

Positions are available at

the University of Aizu

in Japan

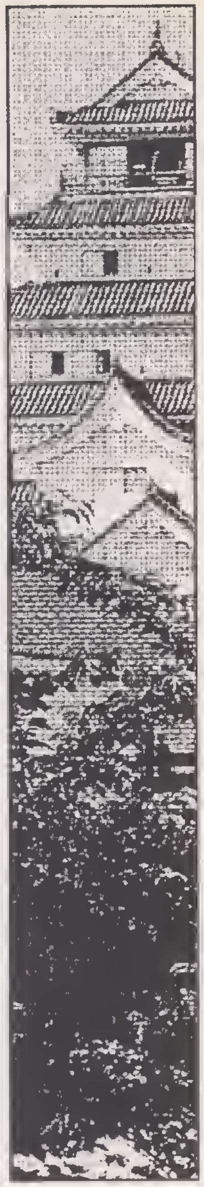

The University of Aizu is expected to be opened in April, 1993. Applications to the faculty positions at all ranks are invited at the Departments of Computer Software, Computer Hardware and Information Systems in the School of Computer Science and Engineering of the University of Aizu. Applicants with research experience in the following areas are of particular interest: data structures and algorithms, distributed and parallel computer architecture, distributed operating systems, compilers for RISC, distributed database management, visual computing, VLSI design, computer networks, sound and speech processing, software architecture and design, performance analysis and modeling, testable logic design, microelectronics, RISC design, computer peripherals, mathematical logic, and differential topology. Each department is scheduled to have positions for 8 full professors, 8 associate professors, 16 assistant professors or research associates. The number of students admitted annually to each department is 80 for under graduates, 40 for masters, 20 for Ph.D. A Ph.D. degree or equivalent is required for a tenure-track faculty appointment, and research associate positions are available for M.S. degree holders. The degrees must be in computer science or a related area. Candidates should demonstrate a strong potential for research and teaching. The salary and rank are commensurate with qualifications and experience. The ability to read and write Japanese is desired but not essential. Serious researchers and teachers of all nationalities are encouraged to apply.

Nominated applicants for the faculty members will be positioned at the university from April, 1993. The positions remain open until filled. The School of Computer Science and Engineering will be staffed with 96 faculties of computer software, computer hardware and information systems, and also with around 20 faculties of liberal arts within the next five years. The facilities include the new buildings worth $300M and computer equipments worth $25M. Research funds will be provided to establish stable research environments. A workstation provided for each researcher will be connected to Internet. Fully air conditioned housing with noise proof flooring will also be provided for the faculty members at the reasonable cost. The University of Aizu, now under construction, is in the scenic vicinity of the Bandai-Asahi National Park, and is located in Aizu-Wakamatsu City, the historical capital of the Aizu region of Fukushima Prefecture, Japan.

Interested persons should forward a letter of application and a complete resume together with the publication list and reprints, at least three letters of reference, graduate transcript, and a statement of your teaching and research philosophy and plans to:

Professor Tosiyasu L. Kunii, Chair
Faculty Search Committee of the University of Aizu
c/o Department of Information Science
The University of Tokyo
7-3-1 Hongo,
Bunkyo-ku, Tokyo 113 Japan
e-mail:kunii@is.s.u-tokyo.ac.jp
fax:81-3-3818-4607

Applications will be accepted until July 1, 1992 at the above address. Late applications are accepted until the positions are filled.

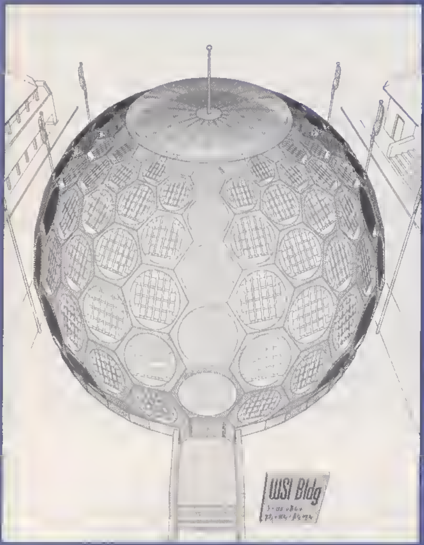

Configuring a Wafer-Scale Two-Dimensional Array of Single-Bit Processors

Ahmed Boubekeur, Jean-Luc Patry, and Gabriele Saucier
National Polytechnic Institute of Grenoble

Jacques Trilhe, SGS-Thomson

This overview of Europe's ELSA project shows how new hardware and software techniques can be used to achieve greater defect tolerance.

Increased integration in digital circuitry can be obtained either by shrinking the transistors, which is ultimately bounded by physical limits, or by increasing the size of the chips up to wafer-scale integration. Researchers have actively studied wafer-scale integration during the past decade and have made some interesting architectural proposals.[1-4] The critical issue is defect tolerance, required to overcome the yield problem; thus, many methods for defect tolerance have been studied (see the sidebar "ELSA configuration compared with other approaches" on page 35). This article describes how defect tolerance is achieved for a wafer-scale architecture implemented in silicon.

The target array is a two-dimensional array of single-bit processors, and the application field is low-level image processing. A processing element (PE) is associated with each pixel of the image. Applications such as image filtering, edge detection, and image enhancement are commonly addressed by this type of architecture.

To achieve defect tolerance, a two-level strategy has been implemented in the ELSA project. The array is divided into subarrays called chips. At the chip level, defect tolerance is provided by an extra column of PEs and bypassing techniques. At the wafer level, a double-rail connection network is used to construct a target array of defect-free chips that is as large and as fast as possible. Its main advantage is being independent of chip defects, as it is controlled from the I/O pads.

Algorithms for configuration play a key role in wafer-scale integration. Three problems must be addressed. The first is the classical problem of constructing an optimized two-dimensional array on a wafer containing a given number of defect-free PEs and connections. The algorithm we describe has proved very efficient. A

0018-9162/92/0400-0029$03.00 © 1992 IEEE

The ELSA project

ELSA (European large SIMD array) is part of the ESPRIT (European Strategic Programme for Research in Information Technology) project 824 on wafer-scale integration, which ran from 1985 to 1991 with a global budget of 10 million Ecus (about $12.5 million). The partners were British Telecom Research Labs in the UK (A. Elliot, P. Ivey, and A. Cox), responsible for the architectural choices and PE design; Darmstadt Technical University in Germany (M. Huch and M. Glesner), responsible for chip redundancy, general power, and clock distribution; the National Polytechnic Institute of Grenoble in France (G. Saucier, A. Boubekeur, and J-L. Patry), responsible for the switching network and reconfiguration algorithms; the Electronic Research Lab of the French Nuclear Center in France (G. Nicolas and C. Massit), for laser fuses/antifuses; the Computer, Robotics, and Microelectronics Laboratory of Monpellier in France, for the floating-gate FET studies; the National Microelectronic Research Center in Ireland (J. Barett), for copper tracking; Thomson Cimsa-Sintra in France (C. Val), for packaging; and SGS-Thomson in France, where project leader J. Trilhe was in charge of manufacturing.

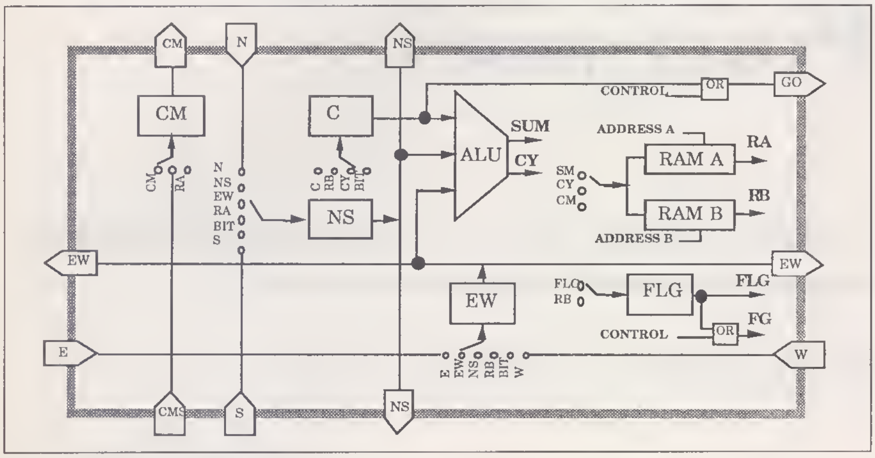

Figure 1. A processing element.

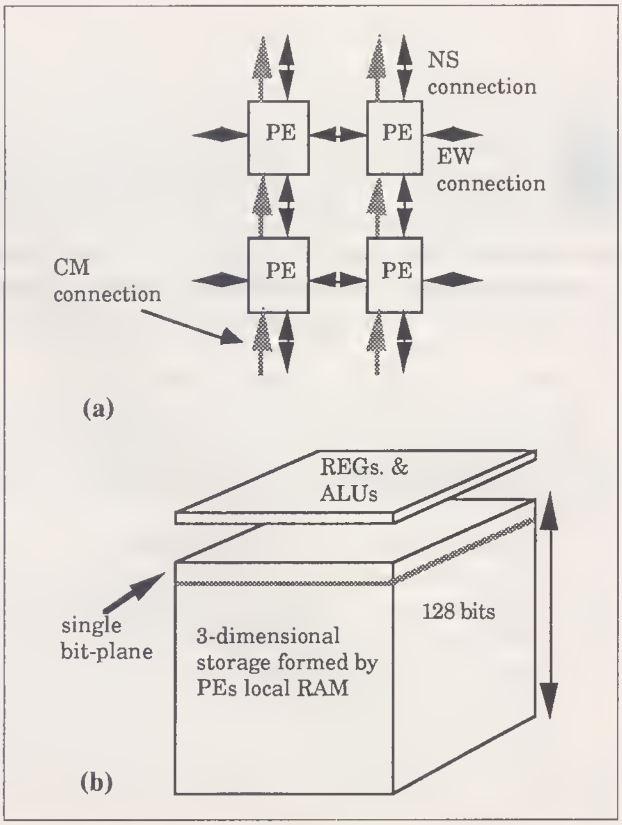

Figure 2. Processing element interconnections (a); a two-dimensional array of PEs with 128 bits of RAM per PE (b).

second problem, more specific to our architecture, is finding a fast way to program the switches for the target architecture found by the first algorithm. Switch programming is done from the border, and the optimization criterion is the reduction of programming time.

The final problem is to program the switches using laser cuts. The programming is done as a final test to verify that the wafer is really operational. The two-dimensional array can still be modified if some presumably good chips and/or switches fail under further testing. After final testing, nonreversible switch programming is performed. A third program minimizes the trajectory of the laser machine.

This article provides an overview of a long-term research effort. It contains the results of the methods and tools implemented during a wafer-scale design, highlighting new approaches. This wafer-scale architecture is called ELSA (see sidebar above).

Wafer architecture and reconfiguration strategy

Target architecture. The target architecture is a two-dimensional array of processing elements operating in a single-instruction, multiple-data-stream mode.

PE structure. The PE is a simple processing cell containing four main functional elements, as shown in Figure 1.

(1) *Storage element:* This element consists of two independently addressable RAMs of 64 × 1 bit each. A multiplexer chooses the RAM input from the

arithmetic logic unit (ALU) outputs or the communication register (CM).

(2) *Communication element:* This element consists of a 1-bit register (CM) that can be connected to its north and south neighbor CM registers. It implements a south-north shift register to scan in (or out) the data from the edge of the array. The CM register can also be connected to the RAM in order to load the data into RAM and get the result out.

(3) *Computation element:* This element contains a 1-bit ALU and its three input registers: north-south, east-west, and C (the carry register). These registers are set from the ALU, the RAM, or other register values. Registers NS and EW can be connected to neighboring PEs for local communication.

(4) *Inhibition element:* A PE can execute either the current instruction or a null operation, depending on a flag register (FLG). The FLG is loaded from local RAM.

The array structure. The PEs are mesh-connected to form an array structure. Each PE communicates directly with its four nearest neighbors, as shown in Figure 2a. In each PE, the local memory can be viewed as a third dimension added to the planar array. Each unit along this third dimension corresponds to a single bit-plane, as depicted in Figure 2b. The array provides additional communication facilities. The arrangement of CM registers, called a plane, is used for I/O data. The NS and EW register planes are used for data processing. A plane of bits in the NS (or EW) register plane may, in one clock cycle, be shifted from one processor to another in one of two possible directions. ELSA has a global input bit in addition to a global output (GO) line. The global input loads scalar data into the NS, EW, and C registers directly. The global output returns global information for the array; GO is pulled low whenever all C registers contain 0's.

The wafer architecture. The aforementioned architecture containing 15 million transistors has been implemented on a 4-inch wafer. There are 20 reticles, each containing four chips, as shown in Figure 3. Each chip contains 12 rows and seven columns, of which one column is a spare (discussed later). These chips are connected through a double-rail switching network. The building block of the switching network is called a switching node, and its function is depicted in Figure 3c. Functionally, all of the switching nodes are identical. However, from an implementation point of view, there are two types: normal and buffered. The signals switched by these nodes are buffered at regular intervals for electrical requirements, to guarantee the integrity of the data. Figure 3b shows the buffered nodes.

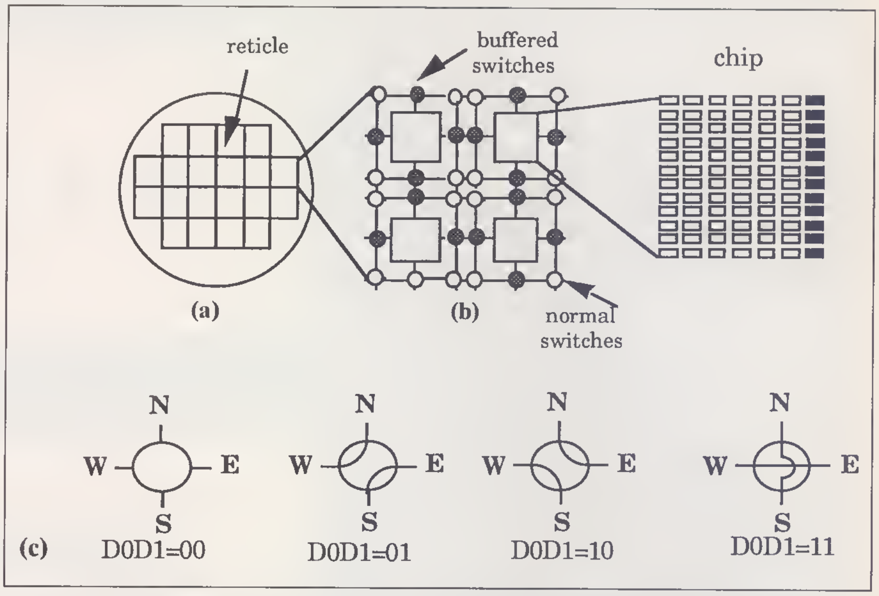

Figure 3. The wafer architecture (a), chips and switching network (b), and switching-node function (c).

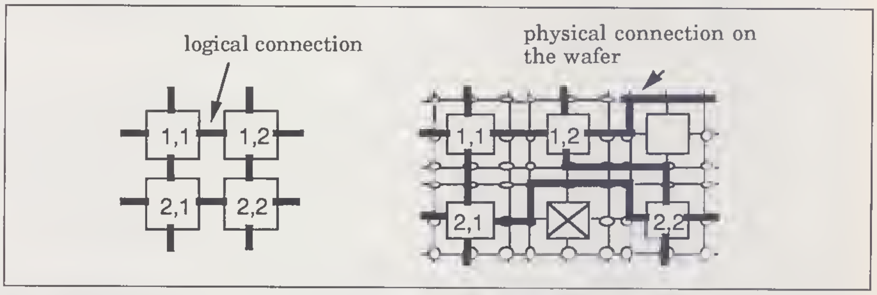

Figure 4. Configuration example for a 2 × 2 array.

Defect tolerance strategy. Defect tolerance is implemented in a hierarchical fashion. First, at the chip level, redundancy has been implemented in the form of a spare column. The chip can be viewed logically as a 12 × 6 subarray. This choice of size resulted from yield calculations and floor-planning considerations.[5] In a row of seven PEs, a faulty PE can be bypassed by transmission gates; hence, six PEs are used per row. Therefore, one PE, at most, can be repaired using transmission gates for east-west and north-south connections.

Second, at the wafer level, the switching network bypasses faulty chips and constructs an operational two-dimensional array with good chips and good connections. This configuration process is illustrated in Figure 4, where a 2 × 2 logical array is constructed from a wafer containing 2 × 3 chips. A connection is a 12-bit-wide data wire plus the control wires required for switching.

The switching network

In the past decade, much research and publishing focused on the technology and structure of switching arrays for wafer-scale integration. Nevertheless, very few were actually implemented on silicon; moreover, the few existing implementations are not realized on a full wafer. Therefore, most represent theoretical work rather than experimental proposals. The sidebar "Switching

Switching networks for wafer-scale arrays

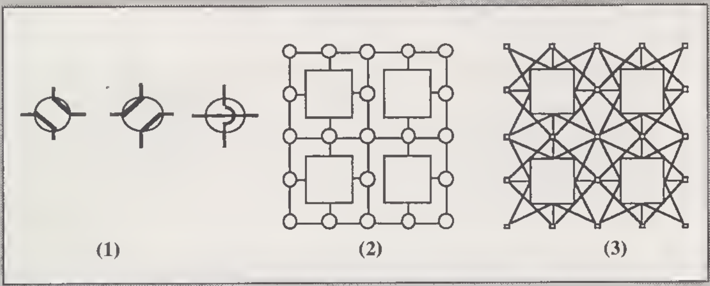

Figure A. A set of switch crossings (1), the mesh array (2), and the star array (3).

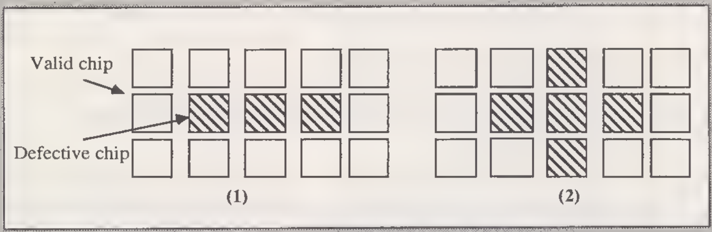

Figure B. A defect on consecutive chips (1); defects in a cluster (2).

A switch is a device that, after manufacturing, establishes or suppresses a connection on the wafer. Switches are classified mainly into two categories: reversible switches, often called soft switches, and permanent or nonreversible switches, called hard switches. The latter imply a modification of the physical structure; a typical example is the laser fuse, wherein a metal line is blown with a laser shot. We suggest the term "firm" switches for those involving a physical but reversible modification. An example is the floating-gate FET, which is programmed with an e-beam but can be erased by either an e-beam or ultraviolet illumination. A detailed classification is available in the literature.[1]

The functionality of a basic switch is defined by a set of crossings. An example is shown in Figure A1.

To understand the functional structure of the array, consider a two-dimensional array of processing elements implemented on a wafer. The switching network aims at constructing, at end of manufacturing, an operational two-dimensional array connecting good processing elements in an way that optimizes performance. The best known arrays in the wafer-scale literature are mesh arrays (Figure A2) and star arrays (Figure A3).

From an architecture point of view, the PEs, or more generally the chips, can be one single-bit PE, as in WASP;[2] an array of single-bit PEs, as in ELSA; or a memory block.[3] The connection can represent *n*-bit buses or a single-bit wire. The choice should be based on the yield simulation according to different chip and switching-array areas.

The defect tolerance of a switching array is of course a key point. Notice that in ELSA defect tolerance is first implemented at chip level through redundant columns, then at wafer level using the switching network.

Defect tolerance must be considered with respect to defects in chips and in the switching array. In WASP[2] and in most hardware implementations, no defect tolerance at all is afforded with respect to defects in the switching network. Implementers more or less justify this through area and yield considerations. On paper, defect tolerance is usually stated with respect to the chips; duplication and triplication are commonly used. Some benchmark patterns could be defined, such as defects on consecutive chips (Figure B1) or on clusters of a given shape (Figure B2), and the robustness of the defect tolerance capabilities of the switching network should be exhibited.

Switch control refers to the device issuing the values controlling the setting of the switches.[4] Such control may be given to PEs or chips controlling their neighboring switches or to external (off-wafer) elements. If a PE contributes to the control of neighboring switches, a defective PE may lead to the loss of all surrounding switches, causing a severe yield loss. This is also true for all switching-array architectures that are supposed to be extended to on-line fault tolerance. The chips themselves are sending messages to route data around the faulty chips.

The quality of a switching network should be asserted in terms of area, performance, and defect tolerance.

- *Area:* If the area overhead due to redundancy is too high, the yield may be affected and careful yield analysis should be carried out. Typically, 10 or 15 percent area overhead is permitted.
- *Performance:* The length of a link between two neighboring PEs or chips defines the array's performance; this is critical for the target architecture. Therefore, buffering techniques are always associated with a switching network.
- *Defect tolerance:* As mentioned above, this is a twofold notion. It concerns first how efficiently faulty chips can be surrounded, and second, how defects in the switching array itself are tolerated.

For ELSA we chose soft switches so that a wafer can be reconfigured if a fault is detected at final test. The definite setting is then fixed with laser fuses to disconnect the device without need of a new configuration. The yield of fuse programming must be very high. We currently achieve a yield of 99.995 percent for programming laser fuses. This has been demonstrated on a test structure composed of 64,000 laser fuses organized as a ROM. On the first wafers, floating-gate FETs have been placed in series as a backup — mainly for research, since the overhead is small (a few micrometers square).

References

1. G. Saucier and J. Trilhe, "Introduction to Wafer-Scale Integration," in *Wafer-Scale Integration*, G.Saucier and J. Trilhe, eds., Elsevier Science, Amsterdam, 1986, pp. 5-20.

2. I.P. Jalowiecki, K.D. Warren, and R.M. Lea, "Wasp: A WSI Associative String Processor," *Proc. Int'l Conf. Wafer-Scale Integration*, IEEE CS Press, Los Alamitos, Calif., Order No. 1901, 1989, pp. 83-93.

3. K. Yamashita, A. Kanasugi, and S. Hijiya, "A Wafer-Scale FFT Processor Featuring a Repeatable Building Block," *Proc. Int'l Conf. Wafer-Scale Integration*, IEEE CS Press, Los Alamitos, Calif., Order No. 1901, 1989, pp. 299-307.

4. M. Blatt, "Yield Evaluation of a Soft Configurable WSI Switch Network," in *Defect and Fault Tolerance in VLSI Systems*, Vol. 2, C.H. Stapper, V.K. Jain, and G. Saucier, eds., Plenum Press, New York, 1990.

networks for wafer-scale arrays" briefly examines the main characteristics of a switching network.

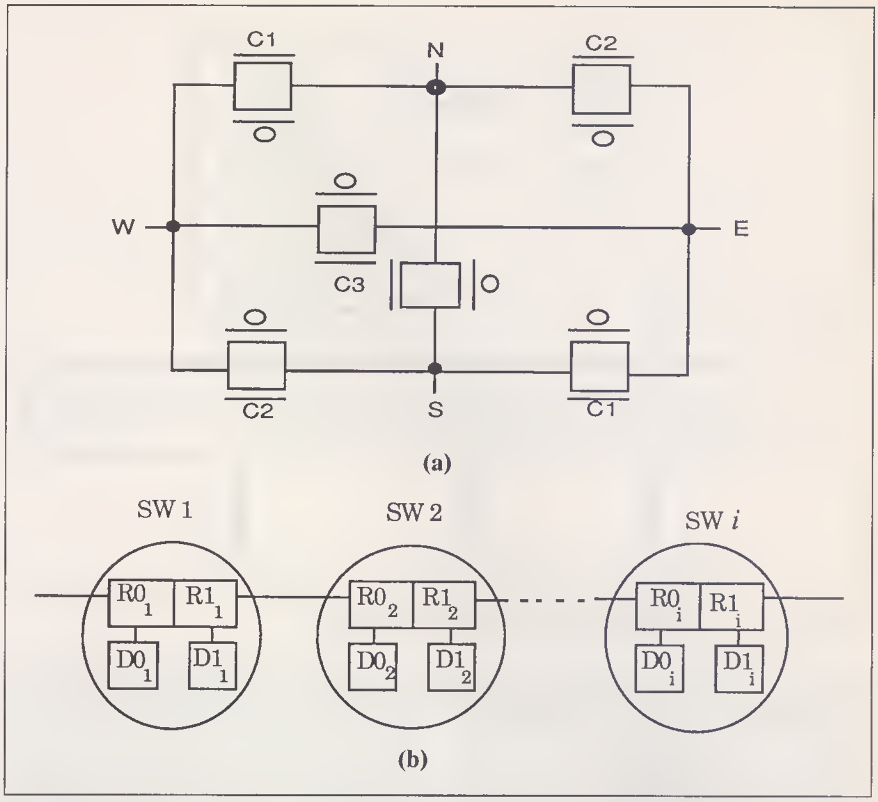

Figure 5. The transfer block (a); programming model of a path (b).

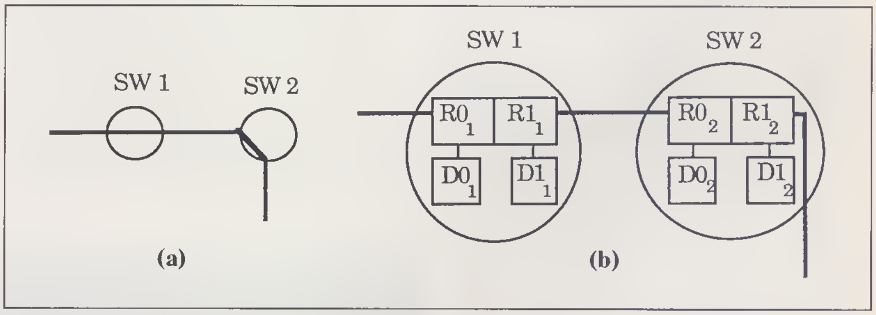

Figure 6. An example of programming two switching nodes: (a) symbolic representation, (b) the model representation.

Basic switch architecture in ELSA. The basic switch has two parts. The first part, shown in Figure 5a, implements the set of paths and is called the *basic transfer block*. The second part is a control part in which two variables stored in two steering latches, D0 and D1, control the transfer block. Figure 3c shows the correspondence between the values of the variables D0 and D1 and the crossing of the switch.

External control of a switching node. To control the state of a switching node, data is sent serially from the I/O pads along a path established between the switching node and an I/O pad. This path can be viewed as a shift register having as input a peripheral switch. A switching node contains two pairs of latches, D0 and D1, controlling the basic switch R0, R1 in which data is stored. Suppose a switching node *i* must be set in a given position through a sequence of switches SW1, SW2, . . . , SW*i*, as illustrated in Figure 5. The data required to place the switches in the desired position is a sequence of values ($D0_1$, $D1_1$), ($D0_2$, $D1_2$), ... This string of values is first pushed using a shifting technique into ($R0_1$, $R1_1$), ($R0_2$, $R1_2$), . . . ; then, at the end, data is latched into ($D0_1$, $D1_1$), ($D0_2$, $D1_2$), ...

Here is how we program a simple path containing only two switching nodes, SW1 and SW2, as shown in Figure 6. To cross SW1 in the east-west direction, the switch needs ($D0_1$, $D1_1$) = (1,1), as Figure 6 shows. So 11 is shifted in $R0_1R1_1$, then this data is stored in $D0_1D1_1$; that is,

$R0_1R1_1 \leftarrow 11$
$D0_1D1_1 \leftarrow R0_1R1_1$

To program SW2, we must cross the first switching node. This is already done. Next, to turn a corner, 10 is shifted into $R0_1R1_1$ while the old value of $R0_1R1_1$ is shifted into $R0_2R1_2$. Then, while the current value is sent to its destination, $R0_2R1_2$, an 11 is sent to program SW1 (see Figure 7).

Switching-node architecture. The switching node has two parts, a control block and a transfer block.[6] Figure 8 shows a block diagram of the switching node.

The transfer block. This block contains 14 identical basic transfer blocks; 12 carry the actual data once the paths have been established. The remaining two (PROG_CTRL and DATA_CTRL) are dedicated to the path-programming phase.

$R0_1R1_1 \leftarrow 10$	$R0_2R1_2 \leftarrow R0_1R1_1$
$R0_1R1_1 \leftarrow 11$	$R0_2R1_2 \leftarrow R0_1R1_1$
$D0_1D1_1 \leftarrow R0_1R1_1$	$D0_2D1_2 \leftarrow R0_2R1_2$

Figure 7. Programming sequence for the example shown in Figure 6.

The control block. This block controls the transfer blocks by feeding them with C1, C2, and C3. The values of these three signals are a function of D0 and D1. The control block holds the programming data in registers R0 and R1, as well as in two other registers, D0 and D1, for storing the D0 and D1 values. Once programming is confirmed, D0 and D1 can be fixed permanently by laser cutting.

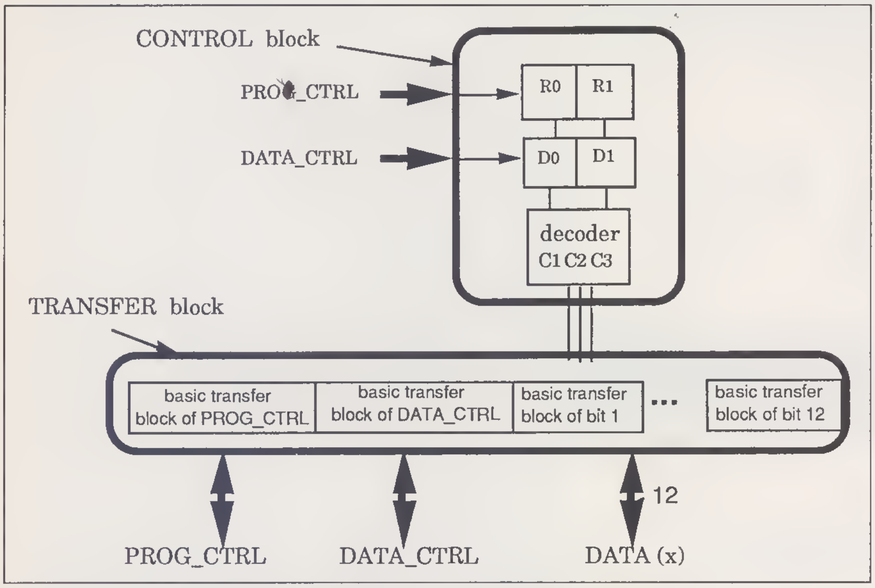
CONTROL block
PROG_CTRL
DATA_CTRL
R0
R1
D0
D1
decoder
C1 C2 C3
TRANSFER block
basic transfer block of PROG_CTRL
basic transfer block of DATA_CTRL
basic transfer block of bit 1
...
basic transfer block of bit 12
12
PROG_CTRL
DATA_CTRL
DATA (x)

Figure 8. The switching-node simplified architecture.

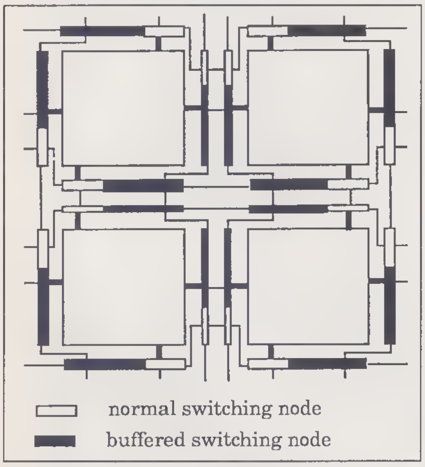

Figure 9. The layout of a chip and surrounding switching nodes.

Layout. ELSA is implemented in 1.2-micrometer CMOS technology. Figure 9 shows the floor plan of a chip surrounded by eight switching nodes. Figure 10 is a photograph of ELSA and its package. Much care has been taken in laying out the switching node to keep its silicon area small. The area devoted to defect tolerance is 18 percent of the total area. This percentage is the ratio of the switching-network area to the total area.[6]

Configuration software

Two algorithms have been developed to configure the wafer. The first constructs the logical target array, optimizing the harvest of good chips. In the first pass, the program generates an initial virtual placement, taking into account the faulty-chip distribution but assuming that all the connections are defect free. The second pass uses a cost function for choosing a chip assignment, taking into account routing capabilities. Routing is done in parallel. The cost function depends on the connection length, the distance between the ideal placement and the current placement, and the deformation of the connection. The second algorithm programs the switching nodes of the target array found by the first algorithm. Recall that this programming is done from I/O pads by sending serial data. Therefore, paths must be found that reach all the switches for programming.

Software defining an optimized target array. An earlier algorithm, called CRAWL1 (for *c*onfiguration and *r*econfiguration *a*t *w*afer *l*evel),[5] enabled us to embed the target array on the fault-free resources of the wafer row by row, column by column, or sector by sector, looking for the best harvest. The disadvantage of this approach was the absence of a global view. It was sensitive to the presence of clustered defects and, more generally, induced a deformation of the target array attracted by the origin of the progressive construction. The choice of the nearest neighbor introduced deformations of logical rows or columns that drastically reduced the probability of achieving a full mapping. Here, we present a new approach that controls the deformation of the target array.

Preprocessing. A preprocessing step establishes a preassignment of the two-dimensional target array. A progressive construction then minimizes the deformation with respect to this preassignment.

After identification of good chips, the problem of preplacement is to find a

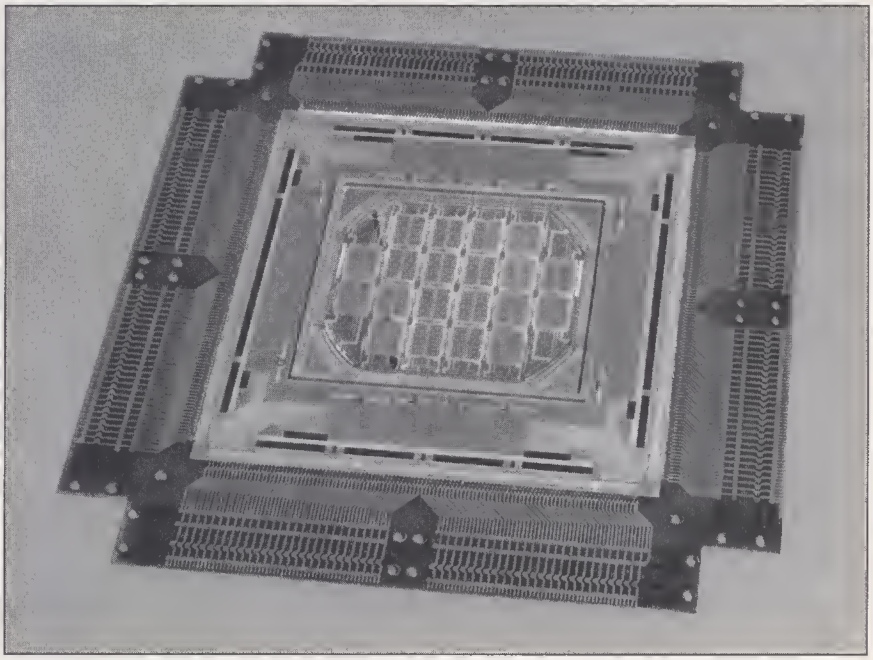

Figure 10. ELSA in its package.

good assignment, assuming the interconnection network is defect free but taking into account the cartography of the wafer. The output of the assignment gives the ideal preplacement with the good chips available on the wafer, assuming all the connections are fault free. For this purpose, we choose a bipartitioning method that always completes. It is based on the divide-and-conquer paradigm and is made up of three phases.

The first phase is recursive. Suppose an $n \times m$ target has been constructed on a wafer containing p valid chips ($p \geq n \times m$). The target is partitioned into two parts (A, B) horizontally or vertically according to the current step of the algorithm (see Figure 11). Let $\lceil x \rceil$ denote the integer part of x.

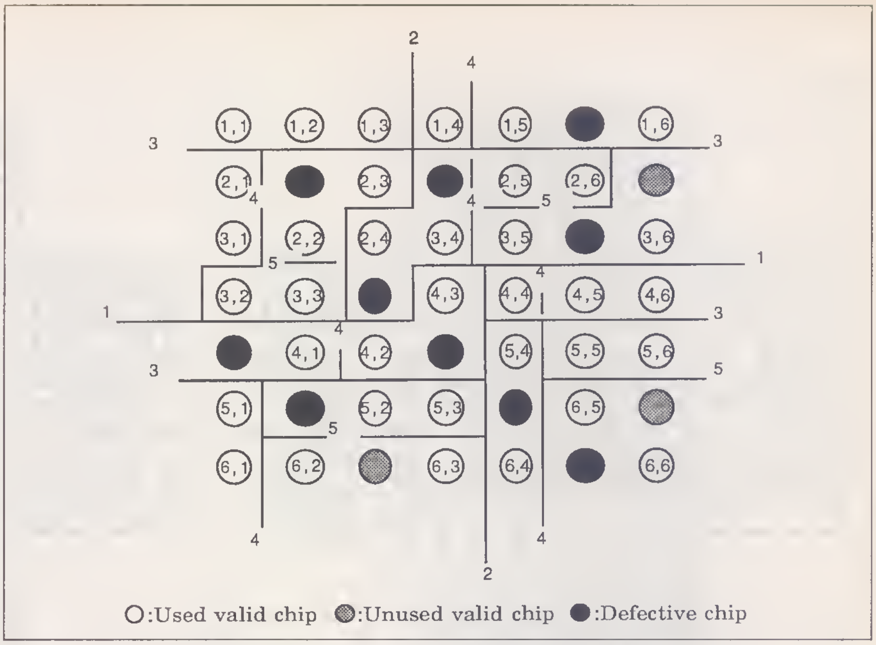

Figure 11. Bipartitioning method on a 7 × 7 defective array.

ELSA configuration compared with other approaches

In the earliest publications,[1] most of the reconfiguration algorithms for two-dimensional arrays ignored defects on connections and switches, and furthermore paid little attention to establishing connections. The theoretical problem consisted of minimizing the longest distance between two neighboring elements. It was assumed that physical elements connecting logical elements of the target were fault free. The distance at that time was more or less abstract, since it was merely proportional to the number of crossed elements. Therefore, this problem was addressed as a mathematical matching problem aimed at associating with each logical element a defect-free physical element[2] with a cost function related to these abstract distances. The difficult issue was the complexity problem. To solve it, hierarchical or divide-and-conquer approaches were adopted. These methods provide a global view of the problem.

A second set of methods[3] relied on a strategy of replacing faulty elements with standby elements. On the wafer, initial columns and rows were identified, and extra columns and rows on the border of the wafer were considered redundant. It was assumed that the "initial" array was as large as the logical target array and that within this array a faulty element would be replaced by a standby element. Within this general framework, the approaches differed by the absence or presence (on one or more of the four sides) of the standby elements, the number of standby rows and columns, and the replacement strategies. If we disregard wiring, the problem can again be viewed as a matching problem consisting of associating each logical defective element with a standby element. In some approaches the replacement strategies are limited, since a faulty element can be replaced by a dedicated standby or from a subset of standby elements. These methods have been called "fault-stealing" methods.[4] The coupling between a faulty element and an accepted standby element is established through a "compensation path."

ELSA differs to some degree from all the other published methods. First, it assumes defects on switches and connections as well as on chips. The resources available on the wafer are considered equivalent, and no distinction is made between standby and initial elements. The goal is to construct the largest two-dimensional array (best harvest) while minimizing the longest distance between two neighboring elements. Therefore, configuration is a better term than reconfiguration.

So stated, the problem is closer to an embedding on a place-and-route problem than a matching or standby-strategy-oriented problem. The approach originally taken consisted of progressively constructing the target array row by row, column by column, or subarray by subarray and concurrently assigning a logical element to a physical element and connecting the logical element through a maze router to the already assigned logical element. This progressive construction suffered from the local view and did not obtain a good distribution of the elements on the wafer. Therefore, the preplacement of CRAWL2 associated with a divide-and-conquer strategy was intended to provide a global view. This phase is closer to the above-mentioned methods, since it assigns a physical fault-free element to the logical element via a bisection method, assuming a good repartition on the wafer. But this is considered only a preprocessing step. The final step definitively places and routes the logical target in a highly efficient manner.

References

1. T. Leighton and C.E. Leiserson, "A Survey of Algorithms for Integrating Wafer-Scale Systolic Arrays," in *Wafer-Scale Integration*, G.Saucier and J. Trilhe, eds., Elsevier Science, Amsterdam, 1986, pp. 177-195.
2. T. Leighton and C.E. Leiserson, "Wafer-Scale Integration of Systolic Arrays," *IEEE Trans. Computers*, Vol. C-34, No. 5, May 1985, pp. 448-461.
3. F. Lombardi et al., "Reconfiguration of VLSI Arrays: A Covering Approach," *Proc. Fault-Tolerant Computing Symp.*, IEEE CS Press, Los Alamitos, Calif., Order No. 778, 1987, pp. 251-256.
4. M. Sami and R. Stefanelli, "Fault Tolerance and Functional Reconfiguration in VLSI Arrays," *Proc. Int'l Symp. Circuits and Systems*, IEEE, New York, 1986, pp. 643-648.

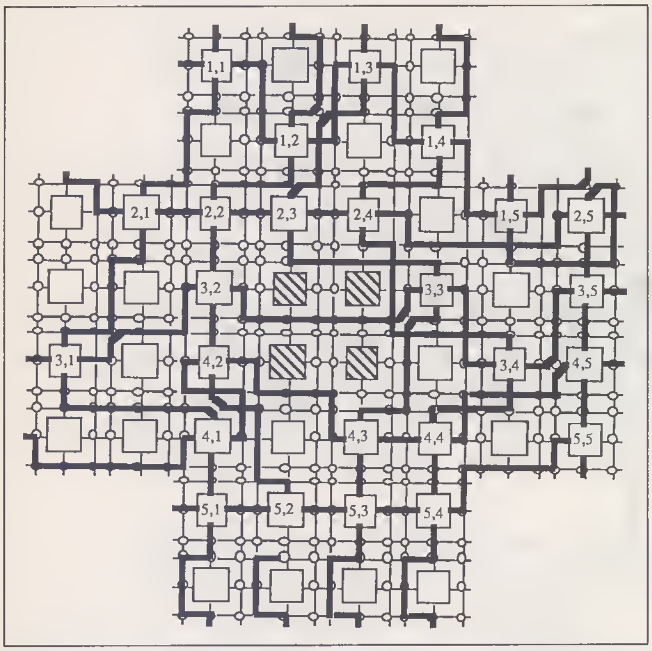

Figure 12. A 5 × 5 array constructed on a cross with a transparent cluster of four defective chips.

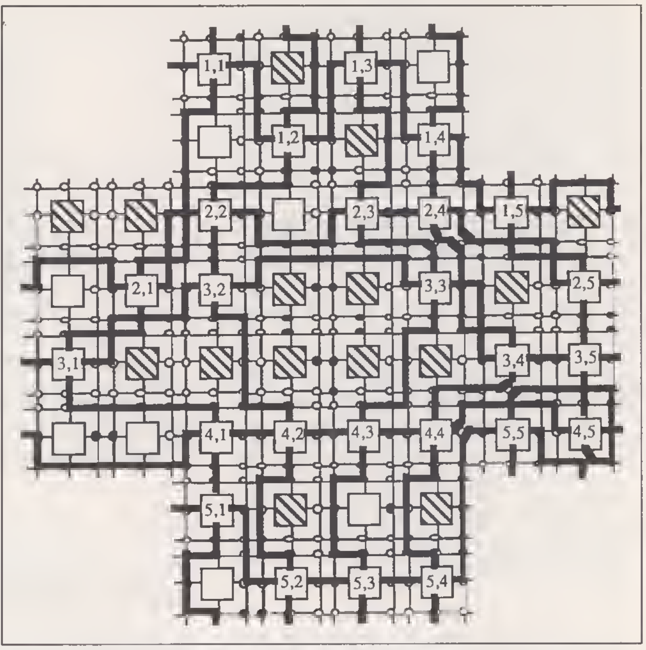

Figure 13. A 5 × 5 array constructed on a cross with an opaque cluster of four defective chips.

- *Horizontal partitioning:* If n is even, the logical array has been partitioned into two $(n/2 \times m)$ subarrays (A, B) containing $p/2$ elements. If n is odd, the logical array is partitioned into two different logical arrays (A, B). The size of A is $m \times \lceil n/2 \rceil$, and the size of B is $m \times (\lceil n/2 \rceil + 1)$. A contains $p \times \lceil n/2 \rceil / n$ elements, and B contains $p \times (\lceil n/2 \rceil + 1)/n$ elements.
- *Vertical partitioning:* If m is even, the logical array is partitioned into two $(n \times m/2)$ subarrays (A, B) containing $p/2$ elements. If m is odd, the logical array is partitioned into two different logical arrays (A, B). The size of A is $n \times \lceil m/2 \rceil$ and the size of B is $n \times (\lceil m/2 \rceil + 1)$. A contains $p \times (\lceil m/2 \rceil / m)$ elements and B contains $p \times (\lceil m/2 \rceil + 1)/m$ elements.

Let's illustrate this using Figure 11, a 7 × 7 array containing 39 valid chips. The horizontal partitioning line, labeled 1, partitions the initial array into two parts. One part contains 19 valid chips and the other 20 chips. The vertical partitioning is labeled line 2 and partitions each half into two subareas containing nine or 10 valid chips. A line denoted 3 partitions these into two subareas of six elements and three elements.

The second phase of the bipartitioning method begins when the current block contains at most two elements. If the block was obtained by a horizontal partitioning, part of a logical row from left to right is embedded. If the block was obtained by a vertical partitioning, part of a logical column from north to south is embedded. An embedding consists of deciding what role a physical element plays in the target array. A chip has physical indices stating its position (row and column number) in the initial physical array. After configuration, physical chips have a row and column number in the target array; these are called logical indices. In Figure 11, the chip having physical indices (3, 2) has become the element (2, 2) of the target array.

The third phase is optional. Since the shape of the successive partitions of the wafer can be irregular, crossings of logical rows or logical columns may occur. To avoid this, the embedding on the frontier row or column (which contains chips of the two areas) is scanned and logical indices may be exchanged.

Cost function. The construction algorithm is based on the evaluation of a cost function that controls the deformation induced by the progressive construction. Each time the next element (in the row or column) must be selected, the element exhibiting the lowest cost is chosen. This cost function F is linear and is expressed as $F = \alpha F1 + \beta F2 + \gamma F3$, where $F1$ is the length of the connection between the candidate chip and the previously placed chip, $F2$ is the Manhattan distance between the candidate chip and the ideal place given by the preplacement, and $F3$ is the distance between the candidate chip and the already neighboring assigned chips.

The parameter $F3$ tries to make a choice consistent with the previous choices to limit the deformation. The array should be constructed from left to right and top to bottom if, for instance, an initial element at the top left of the array has been taken. Any connection to an adjacent element should not go "back" to the left or to the top. Therefore, the coefficient $F3$ tends to place the element (i, j) at the right of element $(i, j - 1)$ and at the bottom of element $(i - 1, j)$.

Practical experiments have allowed us to fix the values of α, β, and γ. If the connection length is limited to k switches, the range of $F1$ is $[1, k]$. For instance, in our real demonstrator, the electrical requirements have limited the path length to 15 switches. To normalize $F1$, $F2$, and $F3$, the value of α is taken to be 1. Practical experience has shown that β must have a strong influence if a small target has to be constructed in a large physical wafer. In this case, the distance to the ideal placement must be severely minimized to ensure a balanced final array. Therefore, for a target close to the largest possible target, β has been taken to be in the range of [3, 5], and for small targets it may be as large as 15. For clusters of defects and nonregular distribution, this coefficient must be increased, and the experimental range has proved to be [5, 10]. These adjustments can be made interactively, or they can be automated. For γ, an average value of 5 appears satisfactory.

Qualitative tests. The evaluation criterion is the harvest (H_n). The normalized harvest is the ratio between the target achieved and the largest rectangular array of good chips on the wafer. Our goal was to produce an algorithm able to overcome some specific manufacturing defects.

Though defect clusters are less frequent in ELSA because of the two-level configuration strategy (the chips are configured using spare columns), they cannot be excluded. Two possible local configurations are tested: transparent clusters and opaque clusters.

- *Transparent clusters:* A group of neighboring chips can be defective even though some of the switches between them remain valid. The CRAWL2 algorithm crosses the cluster; that is, it allows the connections to use the valid switches when they are useful. The example shown in Figure 12 is a 5 × 5 target constructed on a simplified ELSA wafer: A cluster of four chips is crossed by one connection. The normalized harvest obtained by CRAWL2 is 59 percent.
- *Opaque clusters:* This is the worst case, in which an entire part of the wafer becomes unusable. These clusters must be bypassed, and this may increase the length of the connections significantly. Figure 13 illustrates this case. The size of the reached target is 5 × 5, using the same wafer as above. Note that the length of one path is 10 switches (the maximum length allowed is 15). The normalized harvest is 59 percent.

Breakpoint test. Wafers with high defect rates are tested by the breakpoint test. For example, 30 percent and 7.5 percent rates have been considered for chips and switches, respectively. Even with unrealistic defect rates, CRAWL2 produces a configuration with a harvest of 50-60 percent. This can be worthwhile if wafers are to be used despite the poor yield. In the Figure 13 example, for which the defect ratio is 34 percent for the chips and 5 percent for the switches, an array of 5 × 5 is constructed and the harvest is 59 percent.

Quantitative tests. All results hereafter have been obtained with 7 × 7 wafers but can be extended to other sizes. A group containing 50 wafers has been tested. The average rate of defects on these wafers has been kept constant: 10 percent for the chips and 1 percent for the switches. Thus, 50 values have been obtained for each option of the algorithm and average values have been computed. A 77.5 percent average value has been obtained for CRAWL1. For CRAWL2 we obtained 88.7 percent, 85 percent, and 85.5 percent for the column-by-column, row-by-row, and sector-by-sector options, respectively.

Conclusions. First, it has been shown that the previous version, CRAWL1, can be definitively discarded: The addition of an adequate cost function has considerably increased CRAWL2's performance. Second, and this is more surprising, no single method has proved better than the others. The method by row (or by column, which is symmetrical) is more suitable for nonhomogeneous defects as well as for nonsymmetrical targets.

Software for switch programming. The switch-programming phase takes as input the target architecture defined by CRAWL2. The switches and connections of the target array are called useful switches and useful connections. The defective switches, connections, and chips are also located. We will use the target architecture in Figure 14 as an example.

Programming cost. As mentioned earlier, switches are always programmed from the border. Our problem, then, is to program the useful switches of the target array from the pads in minimal time. To do this, the wafer is covered by trees, starting from the pads and reaching all the useful switches. All these trees are programmed in parallel. A cost, defined as the time required to program all the switches on the tree, is associated with each tree. *The optimization criterion of this set of trees is the minimization of the programming time necessary to program the target array.*

Switch-programming problem. The goal is to reach all useful switches while minimizing the programming time. Let S_u be the set of useful switches defined by the target array and D_i be the programming time of the ith tree. The problem is to find a set of n disjoint trees such that the union of the useful switches covered by the n trees is S_u and the programming time D_i of these trees is minimized.

Switch programming by a tree-covering method. The programming will be done in three steps:

(1) Cover the wafer with trees reaching all the valid switches of the wafer. To find the trees,

```
Do forever
  For each pad considered in a
  counterclockwise order,
  construct a tree as follows:
    If there exists valid and free
    neighbor switches for this tree,
    add them to the tree.
  End for
  If no addition has occurred, exit.
Done
```

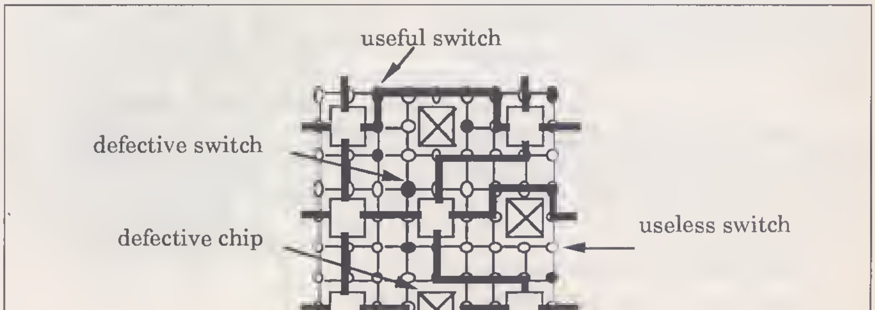

Figure 14. A 3 × 2 logical array on a 3 × 3 defective array.

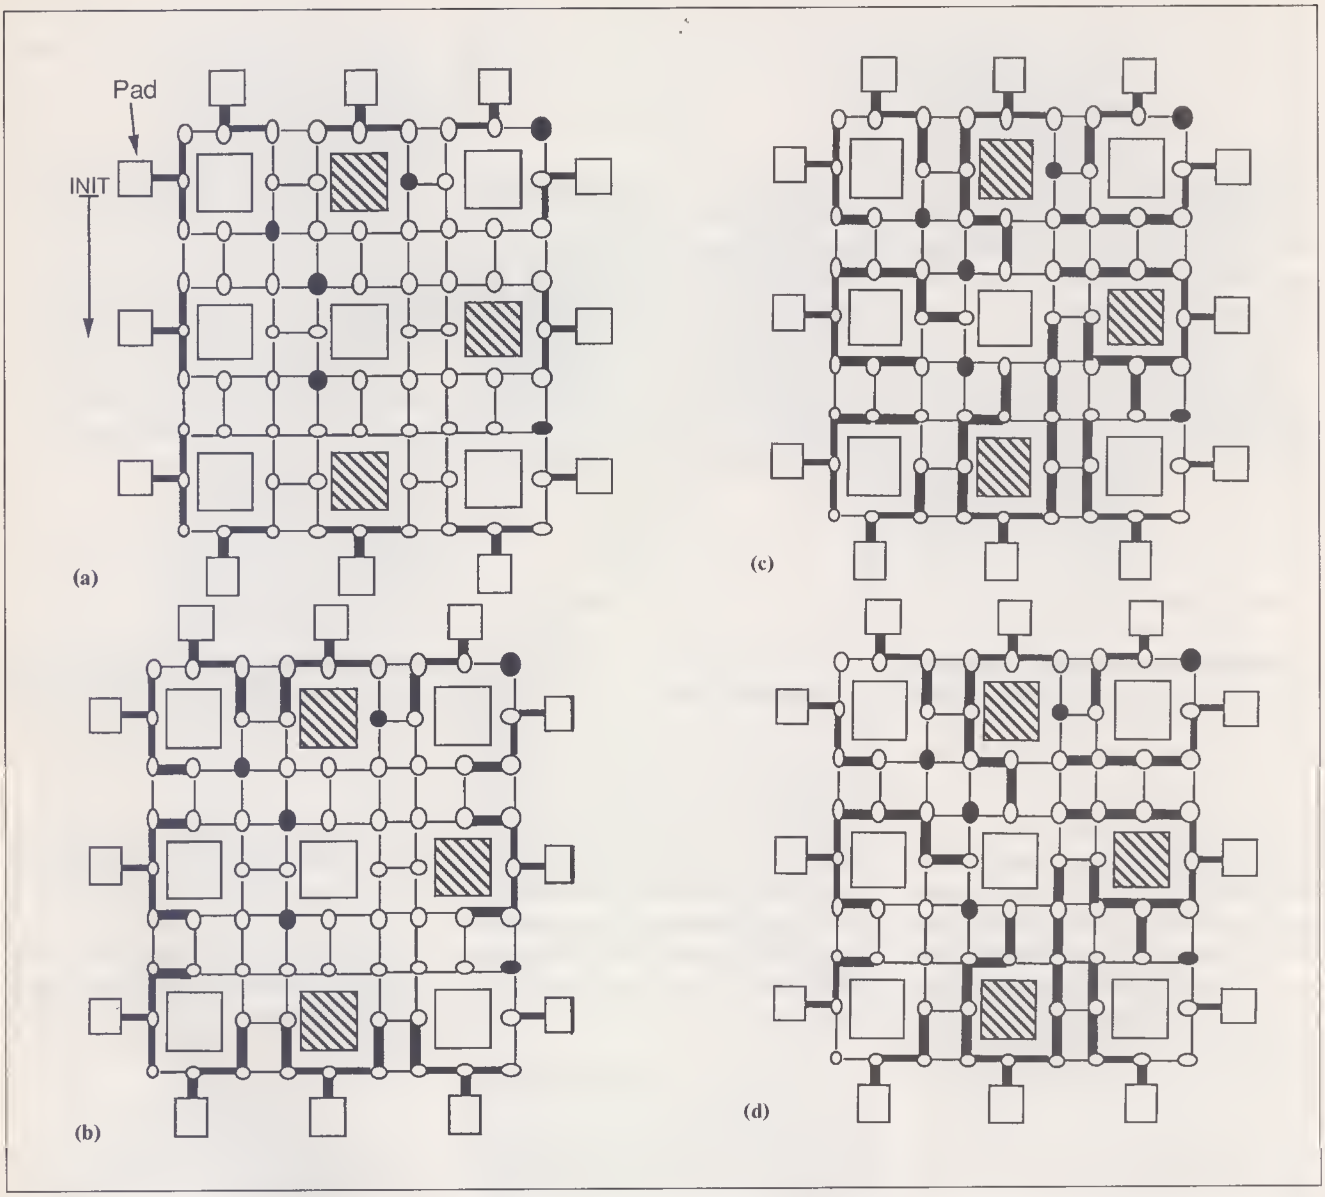

Figure 15. Tree construction: (a) the first step, (b) the second step, (c) the last step, and (d) result of the elimination phase.

Figures 15a and 15b illustrate the first and second steps of the tree construction. The first tree is the one with the pad called INIT. Figure 15c shows the last step of the tree construction; note that all switches of the array are covered.

(2) Reduce the depth of the trees by eliminating the useless switches. Figure 15d shows the result of this step.

```
For each tree
    While the leaves of the tree are
    useless switches, remove them.
End for
```

(3) Carry out the cost optimization phases. Here, the switches that are leaves

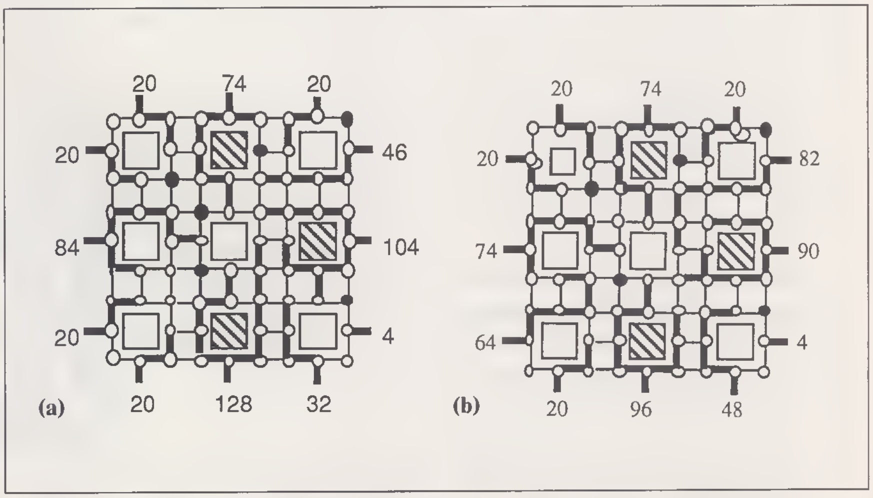

Figure 16. The cost of the trees (a); result of the optimization phase (b).

of deep trees are disconnected and then reconnected to trees of lower cost. Cost optimization is done by the following algorithm:

```
Compute the costs of the trees
Sort the trees according to their
cost
Do forever
   Select the highest cost tree from
   the sorted list
   For each leaf of the tree
      Do forever
         Select the neighbor trees of
         the current leaf.
         Among the trees choose the
         tree (destination tree) that
         would have the smallest cost
         if the current leaf was added
         to it.
         If the new cost of the
         destination tree is smaller
         than the cost of the current
         tree, add the current leaf to
         the destination tree and
         disconnect it from the
         current one. The current
         leaf is set to the "parent"
         of the moved leaf.
      If no addition has occurred,
      exit.
      Done
If there is no decrease in cost, exit.
Done
```

Figure 16 shows the result of this cost optimization phase; the highest tree cost decreases from 128 programming steps to 96. The algorithms have been written in C on a Sun workstation; they are fast and efficient.

The strongest result of this work is the proposal and realization of a switching network controlled externally. Because this switching network induces a reasonable silicon overhead (18 percent) and is not contaminated by defects in the chips, it constitutes a sound solution to a difficult problem. From a software point of view, the proposed algorithm addresses the classical problem of creating a two-dimensional array with high harvest. ■

References

1. P.W. Wyatt and J.I. Raffel, "Restructurable VLSI — A Demonstrated Wafer-Scale Technology," *Proc. Int'l Conf. Wafer-Scale Integration*, IEEE CS Press, Los Alamitos, Calif., Order No. 1901, 1989, pp. 13-20.
2. B. Warren et al., "A 1-Megabit SRAM Fabricated with 1.25 Technology," *Proc. Int'l Conf. Wafer-Scale Integration*, IEEE CS Press, Los Alamitos, Calif., Order No. 1901, 1989, pp. 47-53.
3. J. Trilhe, "A 4-Mbit Static RAM," *Proc. Int'l Conf. Wafer-Scale Integration*, IEEE CS Press, Los Alamitos, Calif., Order No. 1901, 1989, pp. 193-201.
4. K. Yamashita, A. Kanasugi, and S. Hijiya, "A Wafer-Scale FFT Processor Featuring a Repeatable Building Block," *Proc. Int'l Conf. Wafer-Scale Integration*, IEEE CS Press, Los Alamitos, Calif., Order No. 1901, 1989, pp. 299-307.
5. G. Saucier, J-L. Patry, and E-F. Kouka, "Defect Tolerance in a Wafer-Scale Array for Image Processing," *Proc. Int'l Workshop Defect and Fault Tolerance in VLSI Systems*, Univ. of Massachusetts, Amherst, Oct. 1988, pp. 8.2-1 – 8.2-13.
6. J-L. Patry and G. Saucier, "Design of a Universal Switching Network for Reconfigurable 2D Arrays," *Proc. Third IFIP Workshop Wafer-Scale Integration*, organized by IFIP W.G. 10.5, AICA, AEI, Milan, Italy, June 1989, pp. 363-379.

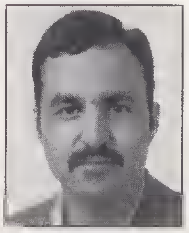

Ahmed Boubekeur is a PhD student at the National Polytechnic Institute of Grenoble, Grenoble, France. His interests include fault- and defect-tolerant architectures, high-level synthesis, and massively parallel architectures for image processing.

Boubekeur received the engineer diploma in electrical engineering from the National Institute of Electricity and Electronics, Algiers, and the MSc from Joseph Fourier University, Grenoble, France, in 1984 and 1988, respectively. He is a member of the IEEE Computer Society.

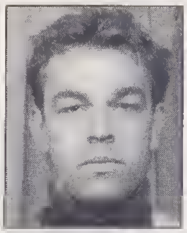

Jean-Luc Patry is a PhD student at the National Polytechnic Institute of Grenoble, Grenoble, France. His interests include fault- and defect-tolerant architectures. He received the MSc from Joseph Fourier University, Grenoble, France, in 1987.

Gabriele Saucier is a professor at the National Polytechnic Institute of Grenoble (INPG) and heads activities at the Integrated Systems Laboratory involving synthesis and test of VLSI, fault- and defect-tolerant architectures, and massively parallel architectures integrated in silicon. Saucier, who received a PhD from INPG in 1970, is a fellow of the IEEE and a member of the IEEE Computer Society.

Jacques Trilhe is the manager of ESPRIT project 824 on wafer-scale integration within SGS-Thomson. He joined Thomson Semiconductors, Grenoble, France, in 1980 to take charge of the pilot line for 2-micrometer MOS technology.

Trilhe graduated from the National Polytechnic Institute of Grenoble with a degree in economic sciences in 1972 and presented his doctoral dissertation in engineering in 1975.

Readers can contact the authors at the National Polytechnic Institute of Grenoble, Integrated Systems Laboratory, 46, Avenue Felix Viallet 38031, Grenoble, France.

PARLE'92

Parallel Architectures and Languages Europe

June 15 - 18, 1992, Paris, France

IEEE Computer Society

ORGANIZED BY **afcet**

CEPIS

COUNCIL OF EUROPEAN PROFESSIONAL INFORMATICS SOCIETIES

The PARLE Conferences have established a worldwide reputation for their excellence and broad coverage. They have set the tradition of a wide and representative international meeting of specialists – from academia and industry – in theory, design and application of parallel computer systems. PARLE has now adopted a new structure and organization to better serve the researchers of the European Community while maintaining its role as the premier international forum for researchers and developers in Parallel Architectures and Languages. The topics of the conference reflect the diverse areas of high quality research and development worldwide.

Sponsorship

The PARLE'92 Conference is supported by the Commission of the European Communities through the ESPRIT Programme and by DGA/DRET, ONERA/CERT, PRC ANM and C3 (CNRS and MRT) ; in cooperation with IEEE CS.

The PARLE'92 Conference is sponsored by CEPIS (Council of European Professional Informatics Societies), EATCS/AFIT (European/French Association for Theoretical Computer Science) and by Bull S. A. (Paris, F), CWI (Amsterdam, NL), GMD (St-Augustin, D), IMAG (Grenoble, F), INRIA (Rennes, F), IRIT (Toulouse, F), LRI (Paris-Sud, F).

Conference Chairpersons

D. Etiemble (LRI, U. de Paris-Sud, F)
B. Lécussan (Sup Aéro,ONERA/CERT, Toulouse, F)
P. Quinton, (IRISA Rennes, F)
J.-C. Syre (Bull S. A., PARLE Steering Committee, F)

Invited lectures

The EDS European Declarative System, C.J. Skelton (ICL, UK)
Scalability Problems in Multiprocessors with Private Caches, M. Dubois (U. South California, USA)
Promises and Issues in Optical Computing, P. Chavel (Institut d'Optique, Orsay, F)
Translating Parallel Programs into VLSI chips, A. Martin (Caltech, USA)
Stanford DASH Multiprocessor : The Hardware and Software Approach, A. Gupta (Stanford U, USA)

Sessions

Architecture
Cache Systems
Data Base Optimization and Modeling
Data Flow Systems
Data Parallelism
Fine Grain Parallelism
Formal Methods
Functional Programming
Graph Theory
Graph Theory and Performance
Implementation of Parallel Languages
Implementation of Parallel Programs
Language Semantics
Network Efficiency
Optimal Algorithms
Parallel Algorithms
Parallel Data Base Machines
Parallel Scheduling in Data Base Systems
Parallel Software Components
Prolog
Semantic models
Systolic Approach

Posters : Tuesday June 16 and Wednesday June 17, 1992 from 14 : 00 to 15 : 30.

Conference location

Bull Center - 5, boulevard Descartes,
Champs-sur-Marne, 77435 Marne-la-Vallée cedex 2

Social events

- **Cruise on the river Seine** : Tuesday June 16 evening.
- **Reception** : Wednesday June 17 evening.

Fees

	Before May 15, 1992	**Beginning May 15, 1992**
CEPIS or IEEE/CS members	FF **2900** incl. VAT 18,6%	FF **3250** incl. VAT 18,6%
Non members	FF **3200** incl. VAT 18,6%	FF **3550** incl. VAT 18,6%
Students	FF **1600** incl. VAT 18,6%	FF **1600** incl. VAT 18,6%

The complete program and registration form will be sent upon request by :

AFCET, 156, boulevard Péreire, 75017 Paris (France) ☎ : + 33 (1) 47 66 24 19 – Fax : + 33 (1) 42 67 93 12

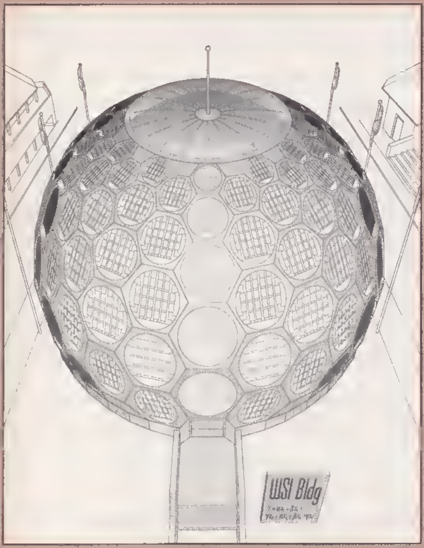

Wafer-Scale Integration Using Restructurable VLSI

Allan H. Anderson, Jack I. Raffel, and Peter W. Wyatt

MIT Lincoln Laboratory

Lasers are used in the final fabrication step for defect avoidance and customization. Nine wafer-scale systems illustrate the technique and its trade-offs.

The restructurable very large scale integration (RVLSI) program at MIT Lincoln Laboratory has established the viability of using a laser to restructure wafer-scale circuits for customization and defect avoidance. Wafer-scale circuits are built with a standard integrated circuit fabrication process when the diffused-link restructuring device is used. (See the sidebar on RVLSI technology.) Wafer-scale implementation requires system, architecture, and technology trade-offs, as demonstrated by nine wafer-scale systems we have built.

To use the RVLSI technique,[1] we first fabricate wafers with redundant circuit modules, called cells, and uncommitted interconnects. After fabrication, we test cells and interconnects, and connect operable cells with a laser to build the desired system. We can also use the laser to customize circuitry and perform diagnostic testing.

As in any integrated circuit, our designs minimize the number of signals at the silicon edge; we design each wafer to be a self-contained subsystem. Our major design decision concerning the subsystem is how to partition it into redundant circuit cells. (See the sidebar on wafer-scale circuit design.) At one extreme, we can partition a wafer so cells comprise a few transistors. Yield will be very high, but area overhead for testing (see the sidebar on testing) and cell-to-cell interconnections will also be high. At the other extreme, if we design the subsystem so cells are large, wafer area will be used redundantly and inefficiently. The trend in our wafer designs has been toward greater cell size and complexity as we moved from a research fabrication facility to a commercial foundry with better yield.

Another important issue is the number of different cell types we should design in the wafer. Both the laws of probability and wafer "floor-planning" considerations favor small numbers. Our primary interest is in signal processing systems, which can be built with many copies of a few cell types. In some cases, we used laser "personalization" to create different circuits from one generic cell.

0018-9162/92/0400-0041$03.00 © 1992 IEEE

Fabricated RVLSI systems

On six different base wafers, we built eight wafer-scale systems and are working on a ninth. Table 1 lists the systems in order of development.

We fabricated the Integrator and Multiplier-Adder Array (MAA) wafers at Lincoln Laboratory on three-inch wafers and the others in foundries using a standard two-micrometer complementary metal-oxide semiconductor (CMOS) process on five-inch wafers. All the base wafers use CMOS circuitry. The "Cells Built" column in Table 1 shows the number of cells of each unique type patterned on the wafer, not including I/O buffers at the periphery. For each type, the "Cells Used" column shows the number of cells required to build the system. The "Total Transistors Used" column shows that the first three wafers had quite small cells.

All systems listed in Table 1 are synchronous. The "Speed" column shows the maximum frequency of their global clocks. All completed wafers have been operated in demonstration equipment. Brief descriptions of the systems follow, illustrating some trade-offs that we made.

Integrator. The Integrator comprises 256 10-bit counters and associated latches, four to a cell. Input data is clocked cell-to-cell at 25 megahertz, but most of the circuit operates at 1 MHz. Latch readout is 1 MHz on a global 10-bit bus. For input data, the system comprises eight linear arrays, and for output, one linear array. One of these wafers has been operating since December 1984 with no vertical-link failures.

MAA. We designed the MAA wafer to implement a completely instantiated 16-point fast Fourier transform (FFT) using 128 serial multiplier-adder circuits and four data-format converters. We customized the wafer using a laser to program the multiplicand and several add/subtract options. We also used copies of the MAA wafer to build a constant false alarm rate filter, a Hough transform filter, and a two-dimensional convolver. All on-wafer communication is serial but the topology is different in each of the four systems. We designed the MAA cells' drive capabilities for the fan-out of two required for the FFT topology. In the other systems, where the input data converter drove more heavily loaded lines, the achievable speed was lower, as Table 1 shows. Partitioning into a "butterfly" cell of four multiplier-adders would let us use area more efficiently but would not let us build the additional designs on the same base wafer.

DTW. The Dynamic Time Warping System is the computation-intensive part of a speech-recognition system. It is a

Restructurable VLSI technology

RVLSI technology uses a laser-programmed antifuse, or link, to make a connection between two metal lines, and laser melting of metal to segment a track into separate nets. The only operational difference between linking and cutting is the laser power level. We developed two antifuses: a vertical metal-to-metal weld and a diffused silicon structure.[1]

For the vertical link, an insulating layer of amorphous silicon or silicon-rich silicon nitride is used between the two aluminum-alloy metal layers. A laser pulse melts both the metal and the insulator to form a silicon-aluminum alloy between the two metal layers. The formed via — viewed with a scanning electron microscope, as shown in Figure A — looks quite different from links formed by short, high-power laser pulses.[2] A connection formed in this way has resistance of less than 1 ohm, so even power-connection links can be small. On the Multiplier-Adder Array and Integrator wafers (presented in the main article) we formed these links with laser linking yields of greater than 99.9 percent. Limited stress tests and operation of an Integrator for seven years show such links to be extremely reliable.

The vertical link's low resistance is a distinct advantage. However, neither amorphous silicon nor silicon-rich silicon nitride is readily available in complementary metal-oxide semiconductor foundries. The diffused link, on the other hand, can be fabricated in any standard metal-oxide semiconductor process. It consists of two diodes separated by a gap equal to the minimum active-area spacing of the integrated-circuit process.

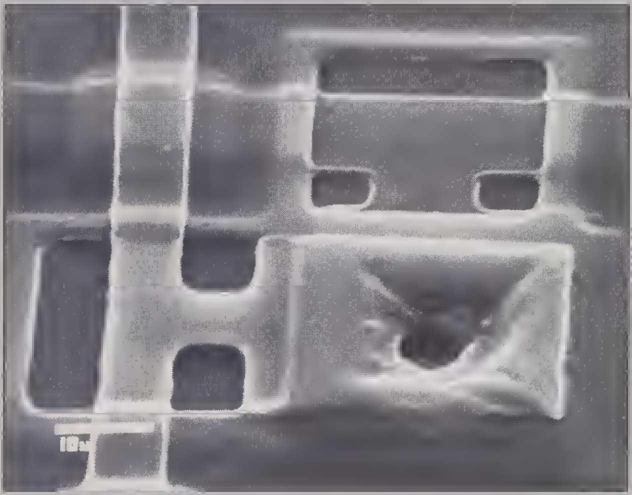

Figure A. Scanning electron microscope photograph of a single vertical link after laser programming.

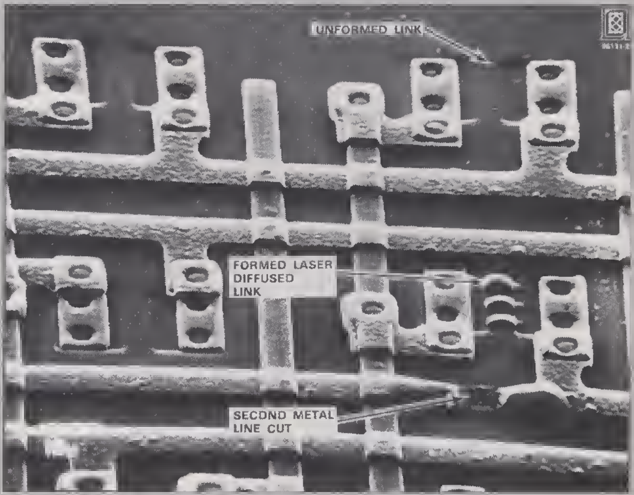

Figure B. Scanning electron microscope photograph of a diffusion link array with tracks on two levels of metal. Cuts can also be made in the first-level metal.

Table 1. Fabricated RVLSI circuits.

System	Wafer	Link	Cells Built	Cells Used	Total Transistors Used (thousands)	Speed (MHz)
Integrator	Integrator	Vertical	96	64	130	25.0
Fast Fourier transform	Multiplier-Adder Array	Vertical	11-352-11	2-128-2	341	10.0
Constant false alarm rate filter	MAA	Vertical	11-352-11	1-144-1	370	8.0
Hough transform filter	MAA	Vertical	11-352-11	2-128-4	352	6.0
2D Convolver	MAA	Vertical	11-352-11	2-156-1	405	3.5
Dynamic Time Warping System	DTW	Diffused	84-84	44-44	374	8.0
Single-instruction, multiple-data	SIMD	Diffused	40-40	25-16	1,884	10.0
Matrix Update Systolic Experiment	MUSE	Diffused	130	96	5,000	10.0
Exponentiator	Exp	Diffused	8-8-4-4	4-3-1-1	500	16.0

The impedance between the diodes is that of a reverse-biased diode, and the capacitance at each terminal is about 30 femtoFarad. A 2-watt pulse from an argon ion laser melts the silicon in the gap, allowing the dopant to diffuse laterally in a few microseconds to create a connection. The merged diode remains isolated from the substrate with a leakage of about 10 picoamperes. The links shown in Figure B are 4 micrometers long and 16 micrometers wide. Linking is done with three laser pulses, and the connection has a resistance of about 50 ohms. Recent designs have 3-micrometer-long links and tighter packing. Depending on glass composition, there may be marking where a link is made, as in Figure B, but there is no rupture.

Power links are very wide. A 1,450-micrometer-wide link, when formed with 240 laser pulses, has a resistance of less than 1 ohm. We can fabricate diffused links with very high yield and detect defective ones when we test interconnections. Laser linking yield is greater than 99.9 percent. We can make cuts in both the first- and second-level metal, with no openings required in either the insulating dielectric or overglass.

The restructuring equipment[1] at Lincoln Laboratory was assembled almost entirely from commercially available components costing about $300,000. The restructuring system positions the packaged wafer under a microscope that focuses the laser beam and lets us observe the operation. When running automatically, the system can make 10 connections or cuts per second, a rate we could increase by pulsing the laser without stopping the table.

Elsewhere, researchers have used several other techniques to do permanent restructuring. Laser-assisted chemical vapor deposition[3] and cutting[4] are somewhat more flexible than our technique. However, like our technique, they require precise mechanical wafer positioning. Moreover, they are slower than our technique and require a toxic gas environment. Fuse-only methods, used in memory chips and wafer-scale-integration memory,[5] do not offer repairability, so they may not work with wafers as complex as ours.

Electrically programmed vias[6] are attractive. This technique is an active research topic in our group, because electrical programming does not require precision machinery and should be faster. However, electrical programming provides no cut capability. Tracks must be fabricated in short segments, so a signal net is likely to have more links in series.

A restructurability technique[7] that uses switches or logic gates to route signals can adapt to in-service failures. Because the overhead is quite large, this technique is best suited to systems with a small amount of interconnection.

References

1. G.H. Chapman, J.M. Canter, and S.S. Cohen, "The Technology of Laser-Formed Interactions for Wafer-Scale Integration," *Proc. Int'l Conf. Wafer-Scale Integration*, IEEE Computer Soc. Press, Los Alamitos, Calif., Order No. 1901 (microfiche only), 1989, pp. 21-29.
2. L. Kuhn et al., "Experimental Study of Laser-Formed Connections for LSI Wafer Personalization," *IEEE J. Solid-State Circuits*, Vol. SC-10, No. 4, Aug. 1975, pp. 219-228.
3. J.G. Black et al., "Rapid Low-Resistance Interconnects by Selective Tungsten Deposition on Laser-Direct-Written Polysilicon," *IEEE Electron Device Letters*, Vol. EDL-7, No. 7, July 1986, pp. 422-424.
4. C.Y. Fu et al., "A Laser Direct Write Double-Level-Metal Technology for Rapid Fabrication," *Proc. IEEE Custom Integrated Circuits Conf.*, IEEE Press, Piscataway, N.J., 1990, pp.18.1.1-18.1.4.
5. R. Bourassa, T. Coffman, and J.E. Brewer, "Progress in WSI SRAM Development," *Proc. Int'l Conf. Wafer-Scale Integration*, IEEE Computer Soc. Press, Los Alamitos, Calif., Order No. 2013, 1990, pp. 13-19.
6. E. Hamdy et al., "Dielectric Based Antifuse for Logic and Memory ICs," *Technical Digest Int'l Electron Devices Meeting*, IEEE, Piscataway, N.J., 1988, pp. 786-789.
7. F. Baba and A. Sinclair, "200-Mbit Wafer-Scale Memory," *Proc. Int'l Conf. Wafer-Scale Integration*, IEEE Computer Soc. Press, Los Alamitos, Calif., Order No. 2013, 1990, pp. 5-12.

Wafer-scale circuit design

Design of an RVLSI system comprises four stages, shown in Figure C:

(1) system and logic design,
(2) physical design,
(3) assignment and routing, and
(4) restructuring and testing.

Physical design is the most complex stage. We must place sufficient quantities of cells and interconnect resources on the wafer, accounting for their projected yields, to construct a system with the desired performance. Wafer length metal tracks for interconnections are segmented later with the laser, so placement of them is not critical; but link placement on these tracks is crucial. Too many links slow the signal paths and reduce track density, while too few makes signal routing difficult.

In practice, we determine the track and link layout incrementally by designing trial layouts, simulating yield, and performing assignment and routing. The bottom dashed path in Figure C shows feedback from assignment and routing experiments to physical design. A layout tool, the Floorplanner, helps us with these experiments. To build more complex RVLSI systems, we must further automate the design phase. An approach we are examining is a library of parameterized link and interconnection patterns to cover most design needs.

Assignment is the process of allocating logical cells required by the system to the physical cells that tested good on the wafer. The layout of interconnection resources constrains assignment, so we can easily perform it manually. In contrast, routing must be done automatically. Even occasional manual intervention at the level of moving routes nearly always causes problems. The Iterative Route and Test tool, which we use for assignment and routing, lets us guide routing by ordering signals, using preferred links, and modifying the cost function used in choosing alternative routes. The tool automates the test-connection generation for optical-probe testing of interconnections and also provides the interface to control the restructuring equipment.

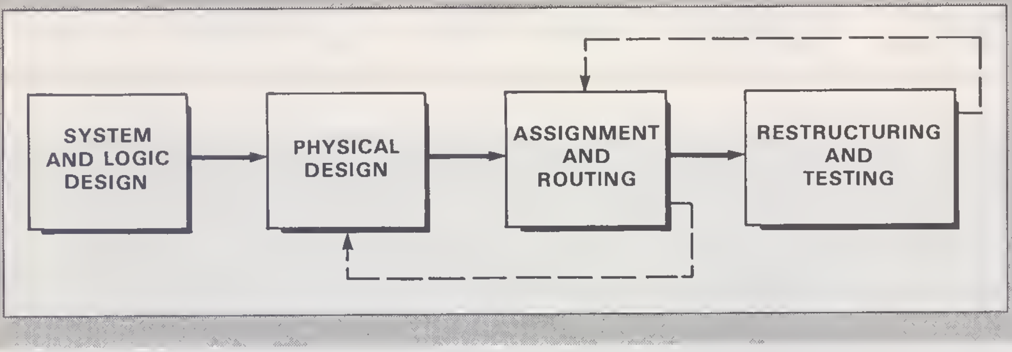

Figure C. Design stages for an RVLSI system. The bottom dashed line shows feedback during the design process; the top dashed line shows new assignment and routing if a cell is found to be bad during restructuring.

linear array of 44 pairs of cells with six global and 13 cell-to-cell signal lines. Of the systems shown in Table 1, it is the only system with active redundancy. Via a serial initialization path, circuitry in each pair of cells can be set to route the cell-to-cell connections around a defective or unneeded cell. During restructuring, we used the bypass feature to test (see the testing sidebar) each pair of cells in isolation.

SIMD. The single-instruction, multiple-data wafer has 40 16-bit processors and 40 2-Kbyte memories. We designed it to support a 4 × 4 processor array imbedded in a 5 × 5 memory array. Although the design is general, we intended it for image processing applications in which an image would be distributed across all memories. Each processor can access its four adjacent memory cells via a local 9-bit bus for seamless processing across image partition boundaries. A single stream of instructions, generated off the wafer, controls the processors and memory. Forty instruction bits are time-multiplexed onto a 10-bit global bus, which is driven from two edges of the wafer for reduced delay and rise time. Because of low circuit yield, we restructured only a 2 × 2 processor array in a 3 × 3 memory array. We designed the cells for a 10-MHz clock rate, which would have required 40-MHz operation on the instruction bus. The integrated-circuit vendor changed the processing of the second-level metal, providing a higher resistance metal and reducing the instruction-bus speed by a factor of 2.

MUSE. The Matrix Update Systolic Experiment is the newest design and a good example of appropriate architecture for the form of wafer-scale integration we describe here.[2] MUSE performs the weight calculation for an adaptive nulling system that combines inputs from a 64-element radar antenna into a weighted sum. MUSE is a linear systolic array of 32 identical modules. The modules' internal circuits operate at 5 MHz. At this clock rate, the system performs 238 million complex number rotations per second or 1.4 billion real operations per second.

MUSE continuously processes data and on external command saves a snapshot of the current state for further processing. An early design had a wide global bus for transferring the snapshot data off the wafer, but this bus and the driving circuits in each cell required a large area. We changed the design so MUSE can perform further processing of the snapshot data in the modules. The results are then passed along the linear array on existing local data connections to an output port. The saving in wafer area offsets the cost of increased control complexity. To save more area, we time-multiplexed the 22-bit data paths onto 11 wires.

Each MUSE module comprises three pipelined Coordinate Rotation Digital Computer (Cordic) circuits that differ only in some control characteristics. Two

Cordic circuits are connected to small memories. We considered three partitioning options:

(1) make the module the replaceable cell,
(2) partition the module into three Cordic units and one or two memory units, or
(3) partition the module into three identical cells, each comprising a Cordic and a memory block.

Our experience with the available foundry led us to conclude that the 145,000-transistor module of the first option would have had too low a yield. The second option had the disadvantages of more cell types and wafer-scale interconnections. With the third option, which we chose, one third of the circuits have unused memories, but the advantages of having one cell type outweigh the area loss. In compensation, the memory and data circuits are powered through separate links so cells with bad memory can be assigned to logical cells that do not use memory. This negates some benefits of having one cell type, because it complicates circuit assignment and leads to longer interconnections.

Building the MUSE system required 96 of the 130 cells on a wafer. Individual cell yield on several multiproject fabrication runs was 75 percent, but on two separate wafer-scale runs fewer than 50 percent of the cells were good. Due to its linear architecture, it was possible to build MUSE on two wafers. A two-wafer system has been built, runs at 10 MHz, and dissipates 4 watts.

Figure 1 shows a packaged MUSE wafer. The 8 × 8-centimeter piece of silicon is cemented at its center to an air-cooled aluminum plate. Wire bonds connect wafer I/O and power pads to the printed circuit board with a number of bonds for power and clocks. Half of the 10 rows of cells are mirrored horizontally for more efficient interconnect-channel use. Six horizontal interconnect channels, including the top and bottom channels, carry power and control signals. The other five horizontal channels are for systolic signal interconnections; each has 85 wafer-length tracks. The only vertical tracks, on the left and right, "turn the corner" for data signals and for fan-out to I/O buffers. Control signals are driven from buffers on both the left and right sides. A one-

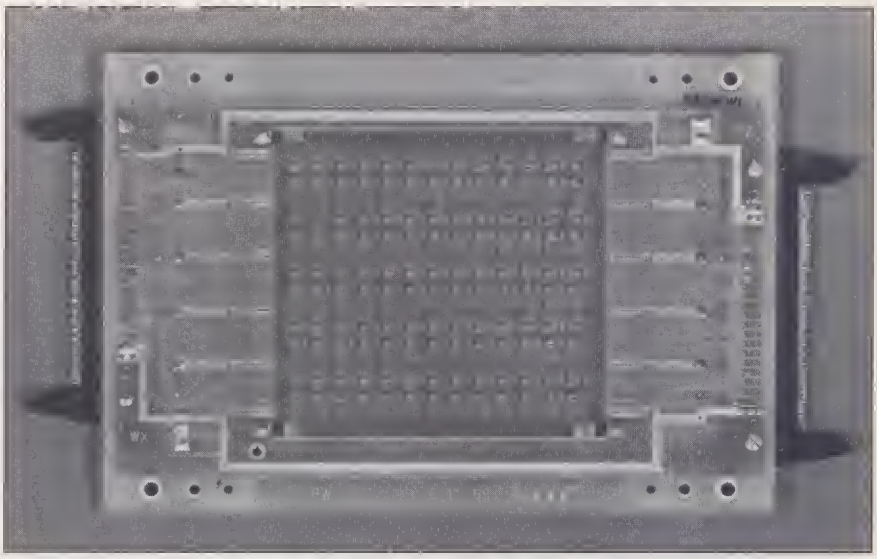

Figure 1. Photograph of a packaged MUSE wafer.

Testing

Because the cost of restructuring each RVLSI wafer is relatively high, the cells and interconnections must be completely tested so only viable wafers are restructured and reworking is minimized. At fabrication time, cells are not connected to the power bus or any other cells so we can test them on a wafer prober. For a high-pin-count cell, such as those in the Exponentiator wafer (presented in the main article), the area required for probe pads can be significant. Therefore, in the Exponentiator we used a boundary-scan path in each cell for functional testing. For measuring track-to-substrate capacitance, all wafer-length tracks and the stubs that connect to cell I/O have a probe pad on at least one end.

We use the measured values and knowledge of the layout to calculate capacitance parameters and a predicted capacitance value for each track. Typically, the procedure is accurate enough to detect tracks with capacitance deviation of about 1 percent from the predicted value. We have seen wafers with all good tracks, but more often wafers have small clusters of defective tracks.

We package the wafer, position it on the laser table, and connect it to the tester before restructuring. A diode to the substrate is located at each input and output of a cell. A special test track, which has links to all signal tracks, is laser connected to a signal track. We measure signal net continuity by exciting appropriate diodes with a low-power laser pulse and measuring the photo current at the test track terminal. As signal nets are tested in sequence, connections to the test track are first made and then cut off. We test partial systems for functionality during restructuring to detect problems as early as possible. All links can be cut off signal tracks so we can replace a cell found to be faulty during restructuring. Our Iterative Route and Test tool (see the sidebar on wafer-scale circuit design) supports such an operation.

New systems should be designed with boundary scan at the cell peripheries to allow better access to the wafer interior.[1] These scan boundaries can also test continuity of interconnections, although not with the accuracy that may be necessary to detect high-resistance links or leaky cuts, and not with the accuracy of optical probing.

Reference

1. D.L. Landis, "A Self-Test Methodology for Restructurable WSI," *Proc. Int'l Conf. Wafer-Scale Integration*, IEEE Computer Soc. Press, Los Alamitos, Calif., Order No. 2013, 1990, pp. 258-264.

Table 2. Area usage on three RVLSI systems.

System	Size (millimeters)	Percentage of Total Area			
		Cells Built	Cells Used	Wafer-Scale Tracks	Power Buses
SIMD	81 × 78	53	29	15	11
MUSE	82 × 88	67	49	13	7
Exponentiator	49 × 49	24	9	47	20

wafer system requires only one connector; the second connector is for a two-wafer system.

Exp. The Exponentiator is very different from the other projects. Nine cells, of four types, are connected to two 64-bit-wide data buses, both of which are connected to package pins; and there are 40 control signals, which come from off-wafer. Due to the large number of I/O pins on these circuits, a boundary-scan register is used for wafer probing. All fabricated wafers have enough good circuits to build a system, but we have not begun restructuring.

Critical issues

Table 2 shows silicon usage for the three most recent designs. Area unaccounted for is area devoted to wafer I/O and simply empty space. The "Cells Used" number is our design target; we did not achieve it with the SIMD and MUSE wafers because of low cell yield.

A single-wafer MUSE system uses 53 percent of the total area for "Cells Used." Of the three systems in Table 2, it was also the easiest for us to design and it was easily restructured. These advantages are due to the linear-systolic architecture of MUSE and its single cell type. It was more difficult to lay out the interconnection resources for the SIMD system due to its two-dimensional mesh topology. We packed the fields of links very tightly to achieve the 15-percent wafer area used by tracks. The "Cells Used" area is only 9 percent of the total area for the Exponentiator for three reasons. The very wide data and control buses occupy nearly half the wafer. Of the four cell types, two are used only once, so statistically the sparing is inefficient. Third, because wafer area was available, we placed more copies of the cells than was necessary.

The suitability of a particular architecture for wafer-scale integration depends on the available technology. The lithography system used determines both circuit performance and the maximum number of cell types patterned on the wafer. Projection lithography, which uses a single mask to print the entire wafer, gives great freedom in logical and physical design. With the electron-beam mask-making used in projection lithography, we can place arbitrary patterns anywhere on the mask. This is important because even a wafer with a small number of circuit types has much nonuniformity at its edges. For one design, we used 33 different pattern files to tile the wafer. However, the relatively large minimum feature size attainable with projection lithography limits performance and the number of cells on the wafer.

To achieve smaller minimum feature size, we can use stepper lithography, which repetitively images a reticle onto a wafer. At Lincoln Laboratory, we demonstrated the production of continuous metal lines at the boundaries between two images, or stitching, with the Integrator and FFT wafers. Other fabricators[3] also use the technique. We placed several images on one reticle and positioned and shuttered it so that only one image printed at a time. But even the MUSE design, which has only one basic cell type, requires two or three reticles with this technique. Electron-beam direct-write lithography preserves freedom of physical design and enables production of high-density circuits. The technique's low throughput precludes its use for high-volume production, but it is compatible with RVLSI methods.

Wafer-scale integration lets us connect circuits with large numbers of pins, but in our designs we reduced interconnections for several reasons:

- antifuse links take space,
- tracks often can't be made as narrow as the technology allows because of their length and high resistance,
- spare lines require additional space, and
- only two levels of metal were available in the foundries used for these designs.

Processes with thicker metal and three or four layers are essential. Thicker insulating layers for reduced interconnection capacitance would also be advantageous. Smaller links[4] would save space and let us reduce the number of spare tracks because a larger population of links makes sparing easier.

The systems we describe here all use CMOS circuitry and run at modest clock rates. But even at these rates, we were concerned about the need to distribute power from the edge of the wafer with only two layers of relatively high resistivity metal. The three most recent systems use gate-dielectric filter capacitors fabricated under the power buses to reduce the voltage drop resulting from peak currents. Higher power circuits will require more metal layers and perhaps power connections to the wafer's interior. Higher dissipation wafers will also require better packaging methods than the ones we used for these circuits. In cooperation with J.E. Brewer and colleagues of Westinghouse Electric Corporation, we operated an MAA wafer in a package which, for mechanical support and heat removal, used a carbon-fiber mounting plate whose thermal-expansion coefficient matched that of silicon.[5]

We developed and demonstrated a technology for monolithic wafer-scale integration through the fabrication of nine wafer-scale systems. The laser interconnection process has high yield and provides the high reliability of monolithic circuitry. The technology is especially well suited to signal processing systems, which characteristically use many replications of a small number of circuits and often have modest interconnection requirements.

Extension of this technology requires integrated-circuit processing with more than two levels of metal and multiple-image stepper or direct-write electron-beam lithography. An improved computer-aided design system could incorporate libraries of parameterized link and interconnection patterns to shorten the physical layout stage of the design process. We envision systems incorporating the best features of hybrid and monolithic approaches. These systems would use chips from other technologies bonded to a wafer to offer such capabilities as optical coupling or high-density memory. ■

Acknowledgments

This work was sponsored by the Defense Department's Advanced Research Projects Agency, the US Air Force, and the US Navy.

References

1. J. Raffel, A.H. Anderson, and G.H. Chapman, "Laser Restructurable Technology and Design," *Wafer-Scale Integration*, E. Swartzlander, ed., Kluwer Academic Publishers, Boston, 1989, pp. 319-363.
2. C.M. Rader, "Wafer-Scale Systolic Array for Adaptive Antenna Processing," *Proc. Int'l Conf. Acoustics, Speech, and Signal Processing*, IEEE Press, Piscataway, N.J., 1988, pp. 2,069-2,071.
3. S.J. Hedge, I.P. Jalowiecki, and R.M. Lea, "WSI Interconnect Issues: Practical Experience Gained on the WASP Project," *Proc. Int'l Conf. Wafer-Scale Integration*, IEEE Computer Soc. Press, Los Alamitos, Calif., Order No. 2126, 1991, pp. 263-269.
4. J. Bernstein, E.F. Gleason, and P.W. Wyatt, "High-Density Metal Cross-Point Laser Linking," *Proc. Int'l Conf. Wafer-Scale Integration*, IEEE Computer Soc. Press, Los Alamitos, Calif., Order No. 2482, 1992, pp. 176-181.
5. R. Bourassa, T. Coffman, and J.E. Brewer, "Progress in WSI SRAM Development," *Proc. Int'l Conf. Wafer-Scale Integration*, IEEE Computer Soc. Press, Los Alamitos, Calif., Order No. 2013, 1990, pp. 13-19.

Allan H. Anderson is a senior staff member at MIT Lincoln Laboratory, where he designs and tests wafer-scale circuits. Previously, he was involved in memory research and the design of high-speed signal processors.

Anderson received a BS in electrical engineering from the University of Wisconsin-Madison in 1957 and an MSEE from MIT in 1959. He is a member of the IEEE and the IEEE Computer Society, and served on the program committee of the IEEE International Solid-State Circuits Conference from 1985 to 1986.

Jack I. Raffel is leader of the Digital Integrated Circuits Group at MIT Lincoln Laboratory. His work has spanned the areas of emitter-coupled logic gate arrays, magnetic film memory, semiconductor memory, analog/digital conversion, integrated-circuit computer-aided design systems, wafer-scale integration, and neural networks.

Raffel received a BA from Columbia College, a BS in electrical engineering from the Columbia School of Engineering, and an MS from MIT, where he was a research assistant in the Digital Computer Laboratory. He is a member of the IEEE Computer Society.

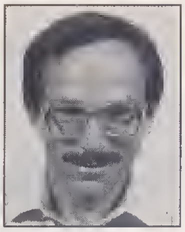

Peter W. Wyatt is assistant leader of the Digital Integrated Circuits Group at MIT Lincoln Laboratory, with overall responsibility for wafer processing in RVLSI and other research projects. His research interests include hardened dielectrics, silicon wafer processing, and wafer-scale integration. Previously, he worked at Bell Laboratories on the development of hybrid integrated circuits.

Wyatt received a BS in electrical engineering from Caltech in 1966 and a PhD in engineering and applied science from Yale University in 1971. He is general chair of the 1993 IEEE International Conference on Wafer-Scale Integration.

Readers can contact the authors at MIT Lincoln Laboratory, 244 Wood St., Lexington, MA 02173-9108. Anderson's e-mail address is anderson@vlsi.ll.mit.edu, Raffel's is raffel@vlsi.ll.mit.edu, and Wyatt's is wyatt@vlsi.ll.mit.edu.

GOING **ONLINE**

The New World Of Multimedia Documentation

SIGDOC'92: The 10th Annual International Conference
October 13 to 16, 1992
The Westin Hotel, Ottawa, Canada

SIGDOC'92 is a conference that will examine design, process, quality, tools, applications, costs and benefits of multimedia documentation.

Presentations will study actual multimedia projects and present research results. Full- and half-day workshops, such as Bill Horton's workshop on visualization, will help you develop your skills. And a participatory demonstration area will let you get your hands on some state-of-the-art online offerings.

Join us as we venture into the new world of multimedia documentation! Register before September 1, 1992, and qualify for early registration prizes.

For more information:
Telephone: (613) 763-2134
Fax: (613) 763-2626
INTERNET address: SIGDOC92@BNR.CA

Sponsored by The Association for Computing Machinery, Special Interest Group on Documentation in cooperation with Northern Telecom and Bell-Northern Research Limited

SIGDOC92 acm

advance announcement

iCSI '92

The Second International Conference on Systems Integration

Headquarters Plaza Hotel • Morristown, New Jersey • June 15-18, 1992

Theme: Managing Large-Scale Integration in the 1990s

This conference focuses on the integration of technologies, processes and systems, and the development of mechanisms and tools enabling solutions to complex multi-disciplinary problems. A special emphasis is placed on the management of large-scale integration. The conference will provide an international and interdisciplinary forum in which researchers and practitioners can share novel research, engineering development, and management experiences. Papers will deal with recent work in theory, design, implementation, utilization and experiences of integrated processes and systems.

Tutorial-in-Brief

Monday, June 15, 1992 • 8:30 am - 5:00 pm

Tutorial 1: Continuous Software Improvement Techniques: The Basis for High Quality Systems Integration
Herb Krasner, Krasner Consulting

Tutorial 2: Eureka Software Factory (ESF)
Bernhard Holtkamp and Herbert Weber,
University of Dortmund

Tutorial 3: The Object-Oriented Database Systems
Wolfgang Klas and Erich J. Neuhold, GMD

Conference-in-Brief

Plenary Session Speakers

- Alan Salisbury, MCC
- David Israel, SDIO/TD
- Richard Kohr, NASA

Technical Program

There will be three days (June 16-18) for 77 papers and eight panels in three tracks. Papers and panels related to the following topics will be presented and discussed:

Issues/Approaches of Integration Track (panels)

- Managing Systems Integration — A Commercial View
- Open Systems
- Integration of Heterogeneous Database Systems
- Managing Systems Integration — A Federal View
- Software Engineering Environment
- Novel Technology Support for Systems Integration
- Integration of People into Large-Scale Systems
- State of Systems Integration Beyond Year 2000

Systems Integration Track (papers)

- Managing Systems Integration
- Systems Integration
- Applications
- Large-Scale Systems Integration in Japan
- Open Systems
- Case Studies

Software Integration Track (papers)

- Software Prototyping
- Requirements Engineering
- Software Process Model
- Software Environments
- Real-Time Systems
- Networking Issues
- Software Development Model
- Databases
- System Specification
- Case Studies

Special Features

- Workshop on Systems Software Automation (Monday night)
- Tutorial on Systems Integration Framework and Architecture (Wednesday Night)

For a complete advance program, please contact:

Professor Peter A. Ng, Institute for Integrated Systems Research, Department of Computer and Information Science, New Jersey Institute of Technology, Newark, New Jersey 07102 USA
Tel: 201/596-3387; FAX: 201/596-5777
E-mail: ng_p@vienna.njit.edu

Sponsored by:
IEEE Computer Society
Institute for Integrated Systems Research, Department of Computer and Information Science, New Jersey Institute of Technology (NJIT)

In cooperation with:
ACM (Association for Computing Machinery)

With corporate support from:
AT&T
Gesellschaft Fuer Mathematik und Datenverarbeitung (GMD)
International Software Systems Inc.

IEEE acm

Conference Committee

Honorary Conference Chair: *Laurence C. Seifert, AT&T*
Conference Chair: *C.V. Ramamoorthy, UC Berkeley*
Program Chair: *Raymond T. Yeh, Intl. Software Systems Inc.*
European Program Co-Chair: *Herbert Weber, University of Dortmund*
Pacific Program Co-Chair: *Fumihiko Kamijo, Tokai University and IPA*
North American Program Co-Chair: *Jules Ghedina, PMG Peat Marwick*
South American Program Co-Chair: *Fuad Sobrinho, EMBRAPA*
Tutorials Chair: *Valdis Berzins, Naval Postgraduate*
Local Arrangement Chair: *Wilhelm Rossak, NJIT*
Steering Committee Chair: *Peter A. Ng, NJIT*

Conference Location

The conference will be held at the Headquarters Plaza Hotel, 3 Headquarters Plaza, Morristown, New Jersey. For reservations call 201/898-9100 or 800/225-1942 or FAX 201/292-0112. Please identify yourself as a Systems Integration Conference attendee to receive the room rate of $96/night single occupancy or $110/night double occupancy. Reservations must be received before May 24, 1992.

ICSI '92 Registration Form • Register Today!

The Second International Conference on Systems Integration, June 15-18, 1992

Please make check, money order, or bank draft payable in U.S. funds to: *Peter Ng - ICSI '92.* Mail form or facsimile with payment to:

Peter A. Ng
Department of Computer and Information Science
New Jersey Institute of Technology
Newark, NJ 07102, U.S.A.
Phone: (201) 596-5749; FAX: (201) 596-5777

Name ______________________________
Last/Family First Middle

ACM/IEEE Membership # ______________________________

Company Name ______________________________

Mailing Address ______________________________

City/State/Zip/Country ______________________________

Work Phone Number ______________________________ FAX Number ______________________________

Do you have any special requirements? ______________________________

Tutorials: Please check the tutorial (full-day) you wish to attend:

☐ Tutorial 1: Continuous Software Improvement Techniques
☐ Tutorial 2: Eureka Software Factory (ESF)
☐ Tutorial 3: The Object-Oriented Database Systems

Conference and Tutorial Registration Fees

(Please check appropriate fee)

Advance Registration Fees - Until May 24, 1992

	Conference Only:	Tutorials Only (full-day):
ACM/IEEE Member	☐ $290	☐ $370
Nonmember	☐ $360	☐ $465
Student	☐ $100	☐ $370

Late/On-Site Registration Fees - After May 24, 1992

	Conference Only:	Tutorials Only (full-day):
ACM/IEEE Member	☐ $350	☐ $445
Nonmember	☐ $435	☐ $555
Student	☐ $120	☐ $445

Total Enclosed (in U.S. dollars) $ ______________________________

If attending both the conference and a tutorial, please add the appropriate fees for your total registration fee. Written requests for refunds must be received at ICSI'92, Department of Computer and Information Science, New Jersey Institute of Technology, Newark, NJ 07102, no later than May 29, 1992. Refunds are subject to a $30 processing fee. Registration fee includes the proceedings and tutorial materials (if attending tutorial). Additional copies of proceedings can be purchased at the conference.

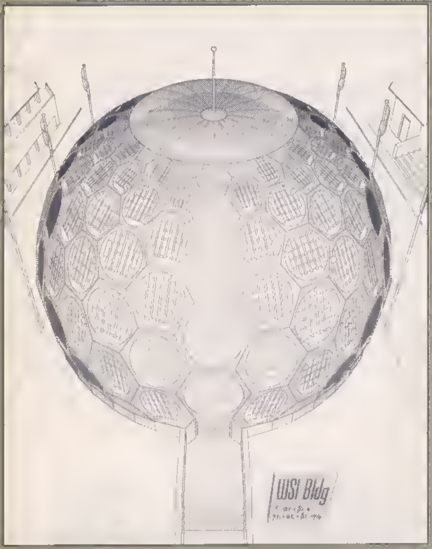

Wafer-Scale Transducer Arrays

Glenn H. Chapman, M. Parameswaran, and Marek Syrzycki

Simon Fraser University

The authors combined laser interconnection techniques and micromachining technology to produce wafer-scale transducer arrays. A prototype thermal-pixel dynamic scene simulator illustrates typical design requirements.

Most current research in wafer-scale integrated circuits has concentrated on building memories, digital filters, multiprocessor systems, or neural networks that are more complex or faster than is possible with the current maximum chip size of about 150 square millimeters.[1] However, large chip size requires defect avoidance; even DRAM memories depend on redundancy and repair schemes to obtain an economic yield.[2] To increase system areas by factors of 5 to 100, we must use wafer-scale integration architectures that divide the target system into repeated circuit blocks (cells) and postfabrication strategies that avoid defective cells and bad signal paths. In such WSI architectures, the structure's larger size is not itself necessarily desirable, but rather a consequence of the limited device density achievable with current fabrication technologies. WSI architecture is a major advantage in applications where the system size, mass, and power are important limiting factors.

Transducers — sensors, emitters, and actuators — are becoming increasingly important in electronics. Microelectronic processing has made the multimicron dimension miniaturization of transducers possible. Microsensors occupy less space and consume less power. Moreover, they influence the observed environment less and offer the potential to integrate complex control and analysis electronics with the transducer.

We speculated that the postfabrication processing technologies used in certain implementations of microtransducers could be combined to create new wafer-scale circuits. In research presented elsewhere in this issue of *Computer*, Anderson, Raffel, and Wyatt demonstrate that wafer-scale integrated circuits can be built using laser-driven interconnection techniques to avoid defective cells after semiconductor fabrication.[3] In other research, micromachining technology has generated several sensors using anisotropic etching of prefabricated complementary metal-oxide semiconductor (CMOS) chips.[4]

We combined these techniques and applied them to applications in which a wafer's extensive physical area is an advantage: systems with large clusters of sensors, emitters, or micromechanical actuators. We call such implementations *wafer-scale transducer arrays*. WSTAs can map the matrix of some variable's distribution within an environment on a very small positional scale for applications such as robotics or process control.

The larger, wafer-scale area of measurement is also an advantage. Emitter WSTAs can inject a controlled matrix of energy (for example, mechanical, thermal, or acoustic energy) into the environment. The energy matrix either acts as a source for other detector arrays or engenders the desired conditions in some medium, again with a very fine positional control. We can combine different sensor

0018-9162/92/0400-0050$03.00 © 1992 IEEE

types within a small physical area, so groups of devices can measure the same physical environment on a small scale locally but track the global changes over a larger region. Emitter-sensor matrix combinations can generate signals and measure the results, with each transducer pair in a localized region. WSTAs also make redundancy strategies more practical for applications in which sensor elements are important for reliability, creating the potential for self-healing or real-time replacement.

Traditionally, sensor arrays — for example, charge-coupled devices — are manufactured using specialized fabrication technologies and large production runs that move along a learning curve to generate perfect devices at reasonable yields. The WSTAs we discuss here are intended for smaller production runs. We use defect-avoidance techniques to achieve the best available transducer-processor matrix for the desired application. Our WSTAs can range in size from moderate areas (a row or column of chips) to full wafer-size devices (25 to 200 times larger than standard chip areas). Applications include temperature-sensor arrays, magnetic-field-sensor arrays (which measure both distribution and multipole moments at many points), smart vision sensors (which do not require every detection site to operate), thermal emitters for creating long-wavelength infrared scene displays, and pressure-sensor arrays (tactile sensors in robotic systems). Some of these applications have been achieved with very large transducers, but WSTAs can tremendously increase positional sensitivity.

In this article we explore the design issues in WSTAs and present a wafer-scale thermal dynamic scene simulator that we implemented using the laser-linking redundancy technique.[3,5]

Transducer production

After a standard CMOS two-level metal process, transducers can undergo

(1) standard manufacturing only (for example, to make photodetectors),

(2) simple unmasked etching (to make thermal pixels),

(3) a single masked etch (to make photodetectors using a back-side etch), or

(4) complex postprocessing, deposition, and definition (to make chemical sensors).

Transducers requiring steps 1 or 2 are generally the easiest to implement as WSTAs.

Many types of semiconductor sensors require no postprocessing and can be fabricated using only standard CMOS or BiCMOS processes. All such sensors measure a nonelectrical variable and convert it into an electrical signal, either current or voltage. Examples are temperature sensors,[6] magnetic-field-sensitive devices,[7] and photodetectors.[8]

Silicon micromachining, an additional processing step, fabricates well-defined three-dimensional structures from monocrystalline silicon and silicon dioxide.[9] Traditionally, fabrication of silicon micromechanical structures for sensor and actuator applications has been performed in customized laboratories. Researchers have explored standard "silicon foundry" CMOS fabrication to develop a more cost-effective approach to creating micromechanical sensors and actuators. This effort has led to maskless postprocessing of CMOS chips to fabricate a variety of transducers[4,10] (see the sidebar on CMOS micromachining).

CMOS-compatible micromachining

Micromachining technology uses the crystallographic etch rate dependency (anisotropic etching) of silicon to produce well-defined three-dimensional micromechanical structures for various transducer applications.[1] Recently, micromachining techniques were introduced as a postprocessing step after commercial CMOS processing to fabricate micromechanical devices for many transducer applications.[2] In CMOS-compatible micromachining, we generate a layout design using a computer-aided design tool and submit it to the CMOS foundry. Then we use a maskless anisotropic etching technique to postprocess chips or wafers produced by the foundry.

The layout design for CMOS-compatible micromachining consists of two parts: the standard design procedure for circuitry supplied by the CMOS foundry and a special derived design that uses standard layout layers but modifies regions on the silicon to form micromechanical structures. The derived design specifies active, contact-cut, via, and glass opening layers, one above the other. CMOS foundries consider overlapping openings as a design-rule violation, but silicon foundries can implement them because the violations do not affect the actual process sequence. The layer definition specifies a portion of the silicon substrate exposed to the ambient environment.

The micromachining process here is a post-CMOS-fabrication etching. The CMOS structures defined by the special layout design are formed by silicon dioxide (field oxide and chemical vapor-deposited oxides) and metal, which are not attacked by the etch. Therefore, no additional masking step is required. Furthermore, the cover glass protects the circuit portion of the CMOS chip or wafer from the etching process. The postprocessing etch is simply a mixture of ethylenediamine pyrocatechol and water.[1] Because standard CMOS devices are implemented on a (100) crystal plane wafer, the anisotropic etch will create on silicon an inverted pyramidal pit bound by the four (111) crystal planes, wherever the substrate silicon is exposed by the layout (see the figure). We can make the silicon openings in any shape to create micromechanical devices for transducer applications.

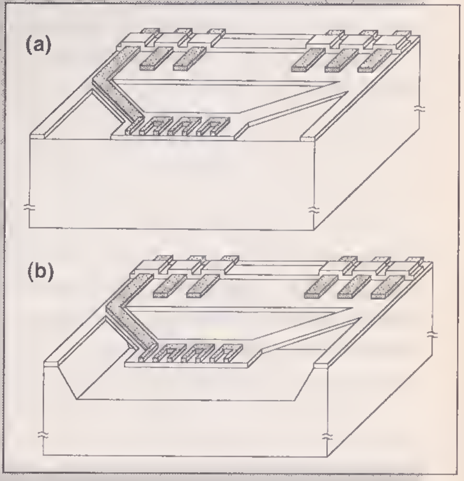

Cross section of an anisotropically etched pit in (100) crystal plane-oriented silicon with a suspended plate of oxide and polysilicon resistor deposited onto it after a standard CMOS process (a) and after anisotropic etching (b).

References

1. K.E. Petersen, "Silicon as a Mechanical Material," *Proc. IEEE*, Vol. 70, 1982, pp. 420-457.

2. M. Parameswaran et al., "A New Approach for the Fabrication of Micromechanical Structures," *Sensors and Actuators*, Vol. 19, 1989, pp. 289-307.

Design factors

The architectures of digital WSI devices and WSTAs share several common points but have a number of differences. Designers of both must break the system into repeated processing blocks, including redundant elements. These are surrounded by communication buses and switches to route signals between the various cells.[3] At the periphery of such devices, there are specialized I/O cells and clock drivers, while the central body contains a mixture of several cell types.

To add a transducer array, we must modify the conventional wafer-scale design. Generally, WSI designs are oriented around digital processing blocks. However, most transducers are inherently analog in part of their function. For example, emitters and actuators are analog, but they can be activated by a digital signal, an analog signal, or a combination of both kinds of signals. Sensors generally output an analog signal. However, transmitting analog signals without losses through signal paths of varying lengths is difficult, and the switching element injects interference. Thus, designs usually use digital processing for data transmission. Such systems often perform considerable analog processing locally but use digital signals for transmission from the transducer cell. For such WSTA systems, a critical design issue is the development of alternatives to many local analog-to-digital converters.

WSTA clock rates are often modest (typically less than 1 MHz) compared with those of digital WSI devices. Hence bus line widths can be smaller and lengths longer: The resistance-capacitance time constant is less important. Serial data transmission of bits (rather than parallel data buses) is often useful at these frequencies to minimize the signal bus area. The exception is devices that dynamically control the transducers' energy emission in real time. Then the question is how rapidly each transducer node must be changed to alter the entire system in the desired manner.

In wafer-scale systems it is desirable to separate the control and processing blocks from the transducer cells both for redundancy (discussed in the next section) and for functionality. Minimizing the number of active devices near the transducer prevents elements of the local environment (for example, adverse temperatures or chemicals) from interfering with circuit operation. Designers should separate active devices from transducers if the specialized processing used to generate the sensors might reduce the fabrication yield of active devices. However, such yield reduction is not an issue with the anisotropic etching method presented in this article.

The large-area device in a WSTA should not disturb the environment. For example, in a temperature-sensor array, we must micromachine isolated silicon sections to prevent heat conduction through the substrate, because silicon is a very good heat conductor.

Transducers are relatively large, so mature, larger geometry fabrication technology can be used to increase circuit yields and reduce costs. However, manufacturer-initiated changes in CMOS technology determine such factors as layer thicknesses and polysilicon sheet resistance. The consequent alterations in transducer design to ensure proper operation may reduce sensor efficiency.

Redundancy in WSTAs

The key to any successful wafer-scale system is a procedure to substitute spare cells for failed circuit elements.[3] With regular WSI devices, the physical position of a substituted cell is restricted by the ability to reroute signals from the inoperable circuit block to a spare block and by limits on the new path's delay time. With WSTAs, designers operate under an additional constraint. The granularity in location (the allowed variation in the physical position of a substituted transducer) depends on the application, but generally it is very limited. For sensor arrays, global redundancy is effective only in cases in which delivery of the same environment at each cell site is possible. Thus, sparing is local rather than global: The substitute cell must be a nearest neighbor of the failed one rather than a cell anywhere on the wafer. Because the processing and control electronics block has a nearly global sparing capability, it should be separate from the transducer cell.

A problem with any local sparing redundancy scheme is that most fabricated wafers suffer from cluster defects: Faults tend to be adjacent. Indeed, Stewart reports that as wafer yield rises, the remaining defects tend to cluster.[2] This phenomenon limits the usefulness of digital WSI redundancy schemes that rely only on nearest neighbor substitutions. However, inherent advantages in WSTAs reduce this effect and allow architectural procedures that help compensate for it.

In comparison with WSI cells, the transducer's greater size reduces the probability of a defect that would eliminate a smaller geometry transistor, also destroying the sensor. In addition, the area cost of designing multiple vias or contacts, which are common failure points, is very low relative to the overall transducer size. Transducer structures are also often less complex and involve fewer layers than circuit blocks. Thus, they are less susceptible to failure after they are built. Indeed, because the size of micromechanical structures is often fixed by physical operation requirements, advances in processing smaller devices will enhance the yield of transducers. However, some transducers require unusual patterns, special processing steps, and fabrication control that may interfere with the yield of the regular electronics.

In WSI digital devices, the cell can be verified in principle with the proper set of test vectors. However, often we cannot fully test transducers by simple electronic means. Instead we need an external controllable source of the desired variable (for sensors) or measurement of the variable generated (for emitters or actuators). In addition, we may need to program calibration factors for each transducer. For certain transducers, we can provide some level of self-test to indicate a few failure conditions (for example, transducers that check their internal resistances). We must also solve the ever-present problem of power connection. Because a high percentage of failed cells on a normal wafer contain power shorts, only good sections can be powered up.

Wafer-scale defect-avoidance methods that involve little or no postfabrication processing on standard CMOS wafers are clearly preferable to methods requiring special masking or fabrication steps. Active switching elements have been widely discussed but are difficult to implement successfully as the only redundancy method on WSI systems. The seemingly low-level problems of shorted or unpowered cells (resulting in the signaling of unpowered devices) and bad signal buses (broken and

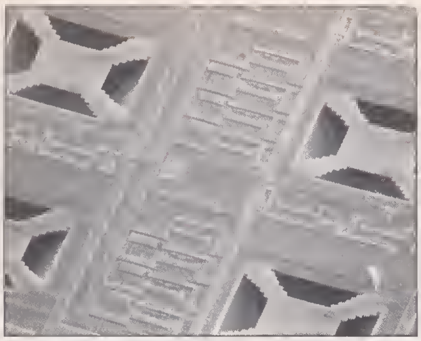

Figure 1. Micromachined thermal-pixel array with adjacent control circuitry. The oxide plate is suspended above the etch pit generated by the anisotropic etching. A polysilicon resistor sits on the plate and is connected to the current source via metal lines that cross the bridges from the plate to the substrate.

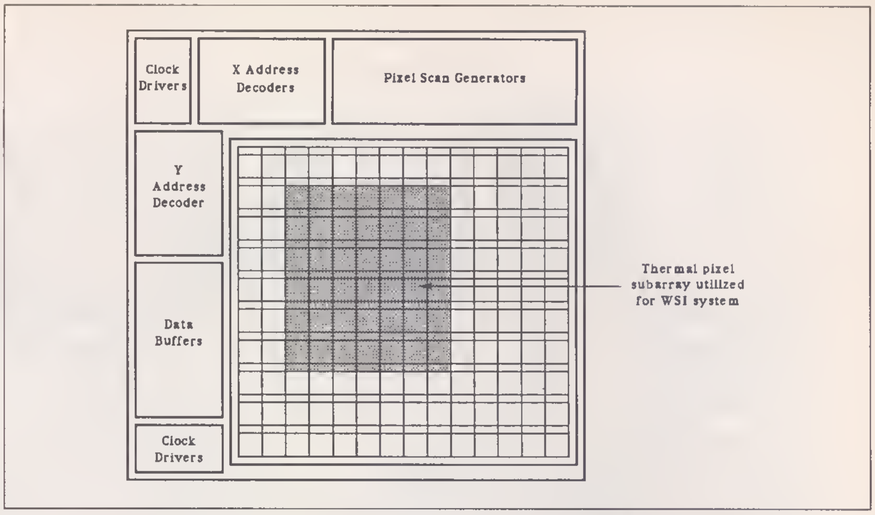

Figure 2. Block architecture and an example location of a thermal-pixel scene subarray corresponding to a specific defect distribution. The emission scene can be moved around to compensate for defects that develop over time.

shorted lines) create substantial complications.

A good first approach to defect avoidance for WSTAs uses a "hard-wiring" method, such as Lincoln Laboratory's laser-linking process.[3] Designers use a laser's power to make physical, low-resistance connections between metal bus lines to segment the bus paths and to route signals to working devices and avoid defective ones. For WSTAs, we use a fully CMOS-compatible structure known as the laser-diffused link.[3,5] It consists of two lines of conductively doped silicon separated by a narrow gap. When selectively exposed to the beam pulse of an argon ion laser, the dopant spreads throughout the gap, reducing the resistance between the lines from that of two reverse-biased diodes to about 100 ohms (depending on wafer dopant levels).[3,5] Longer dopant sections produce lower resistance values for connecting the cell's power. The links have a very high yield (the failure rate is less than 0.01 percent). Laser linking is probably the most successful WSI procedure: Anderson, Raffel, and Wyatt have successfully implemented five different designs, as they report in their article in this issue of *Computer*.[3]

We use the laser-linking procedure within the array to make each transducer from a small matrix of subcells, called subtransducers, several of which are combined to form the full transducer. After fabrication and testing, we "harvest" all the working devices using laser-linking connections to hook them up to the power and ground, the required processing elements, and the signal buses.

To obtain a level of global redundancy, we designed a transducer array system that involves considerably fewer elements than the physical number of rows and columns in the wafer-scale matrix. By selectively activating the working cells, we can adjust the position and size of the transducer array. We create an "origin shift error correction" by shifting the rows and columns of the working section. This design lets us avoid defective areas (most common at the outer wafer edges) and provides self-healing and repair capabilities through subcell substitution.

Thermal-pixel dynamic scene simulator

To explore the implementation of WSTAs, we chose a thermal-pixel dynamic scene simulator as a learning vehicle. Calibration of the complex infrared optical systems used in research, medicine, and military applications requires continuously changing two-dimensional thermal scenes composed of submillimeter-size thermal spots. Conventional thermal resolution test targets and dynamic scene simulators are large and lack resolution.[11] We fabricated miniature heater elements using microelectronic processing and silicon micromachining techniques,[10] and integrated them with digital scene simulator support circuitry.

Many proposals for integrated thermal dynamic scene simulators require customized fabrication steps involving expensive, specially designed facilities. Our alternative process uses a CMOS-compatible thermal picture element. As Figure 1 shows, the pixel is a thermally isolated heater filament that can generate a local hot spot. The small blackbody emitter produces elevated temperatures, emitting radiation in the infrared submillimeter wavelength. In addition, digital control of the heater's current produces a range of infrared intensities for a thermal gray scale.

CMOS technology provides the thermal pixel as a polysilicon heater filament sandwiched between the field oxide and the chemical vapor-deposited oxides. After wafer fabrication, a single anisotropic etch of the silicon substrate produces a cavity for thermal isolation in the silicon underneath the polysilicon heater. The heat isolation is sufficient for a pixel to glow incandescently with about 3.4 milliamperes of current for a short time. For the target system, the desired temperatures are much more moderate (100° to 200° C) than the temperatures accommodated by the heat isolation cavity, and device lifetimes appear to be extensive.[10] Current postetching yields on multiproject test chips are about 65 percent. The intensity of the emitted radiation follows a blackbody relationship and is proportional to a T^4 relationship (T in degrees Kelvin). Simple heat-flow calculations for the isolated polysilicon resistor show that T is proportional to the resistive power loss I^2R.

The wafer-scale thermal-pixel dynamic scene simulator contains an array of thermal pixels emitting infrared energy and control circuitry to simulate a thermal scene. Figure 2 shows the block architecture. The system operates like a

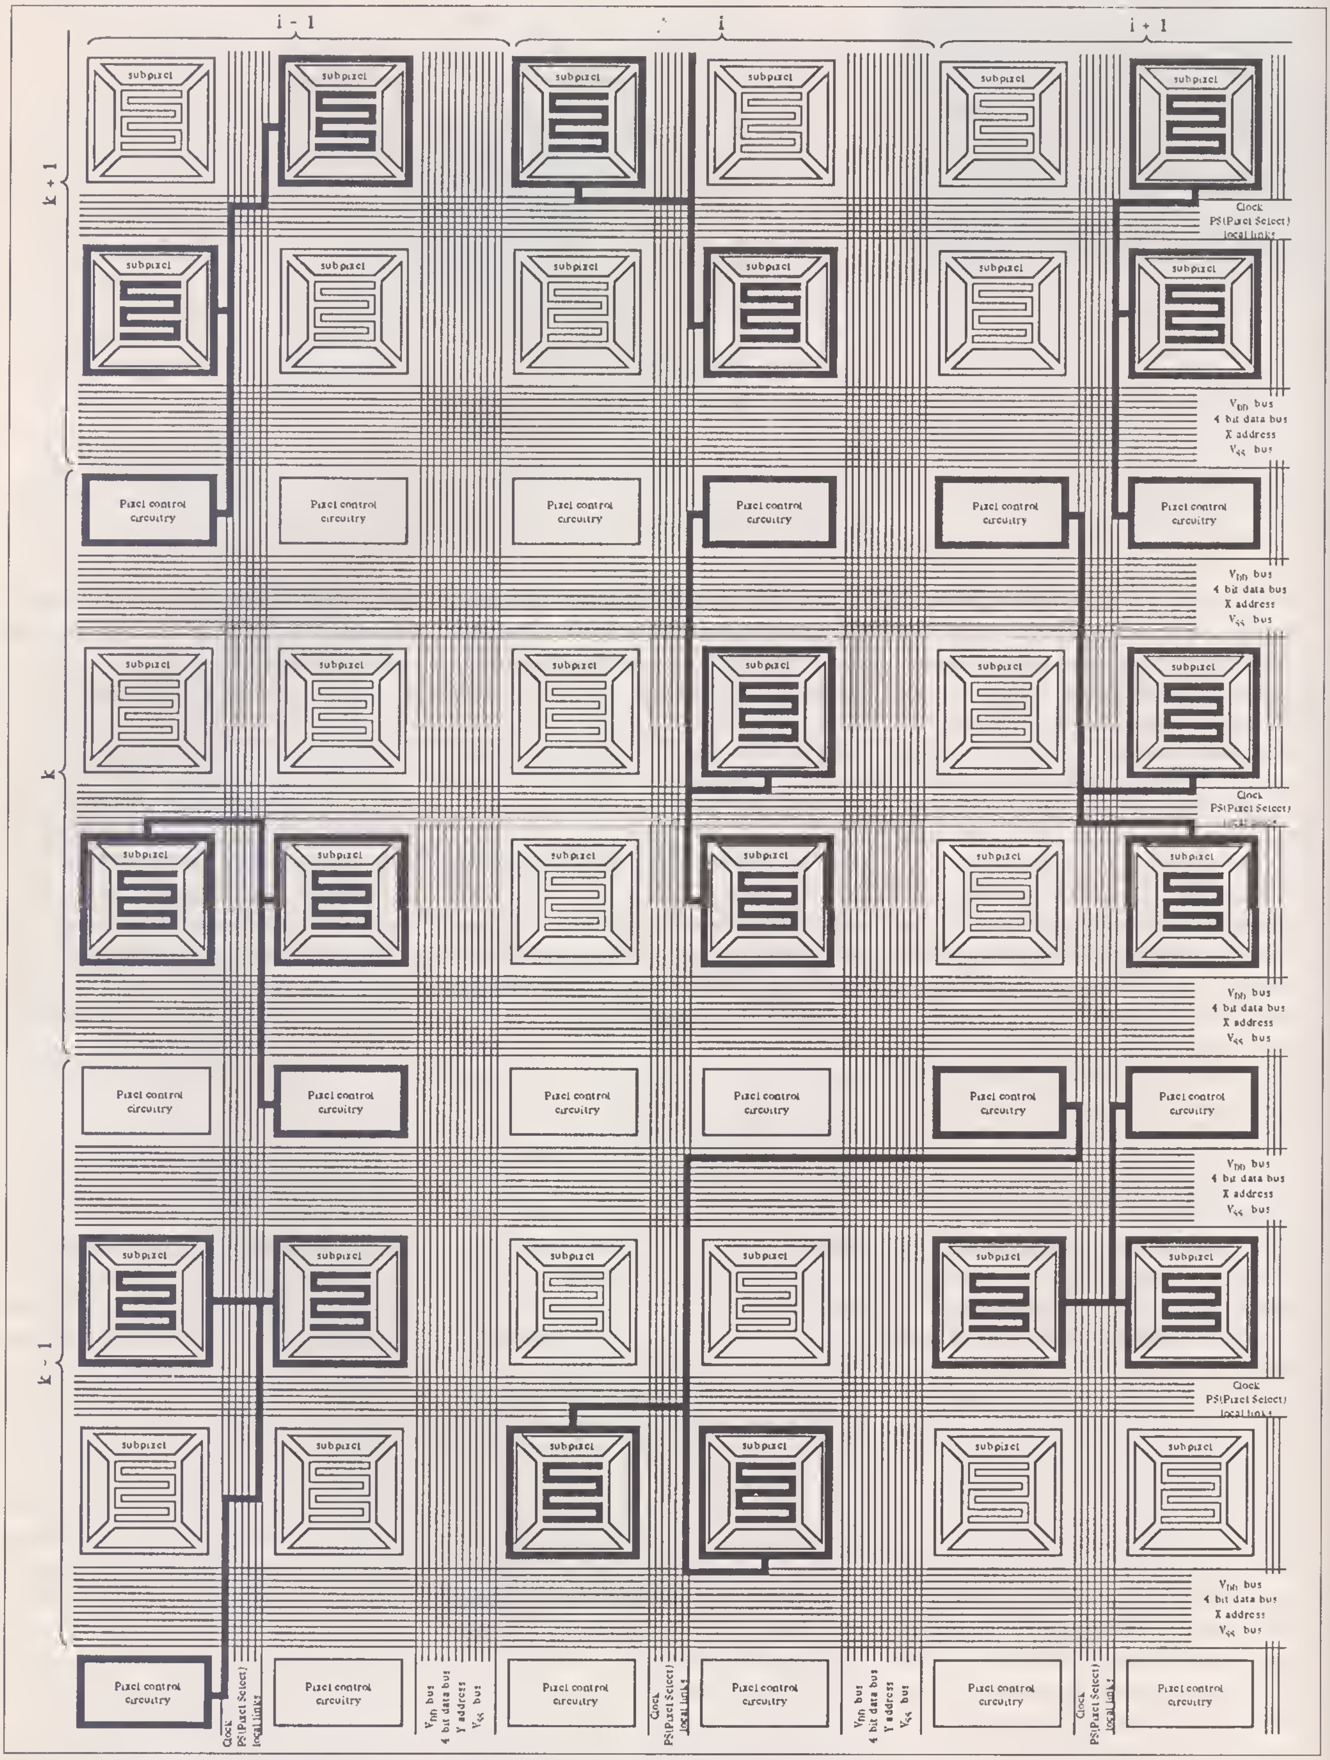

Figure 3. Examples of sparing transducers and control circuits within a wafer-scale thermal-pixel array. Although the shapes of full pixels can change, pixels are confined to specific row and column pairs.

memory-mapped display, with each pixel containing an intensity register. The periphery of the wafer contains *x* and *y* physical-addressing decoders, clock drivers, and data bus buffers. The pixel intensity data is loaded in parallel into the data buffers, which place it on a wafer-wide data bus, where the addressed element latches it into the thermal pixel's display register. This data represents the magnitude of a current flowing in each pixel and thereby an intensity of thermal energy radiated from each pixel to the surrounding environment. The present design uses 4-bit intensity data to program 16 gray levels of thermal energy. A pixel scan generator controls the thermal scene displayed by the array and drives an on/off switch for all pixels. Clock drivers synchronize dataflow to the registers associated with each pixel. Multiple copies of these blocks are built in, providing the global redundancy necessary for WSI devices.

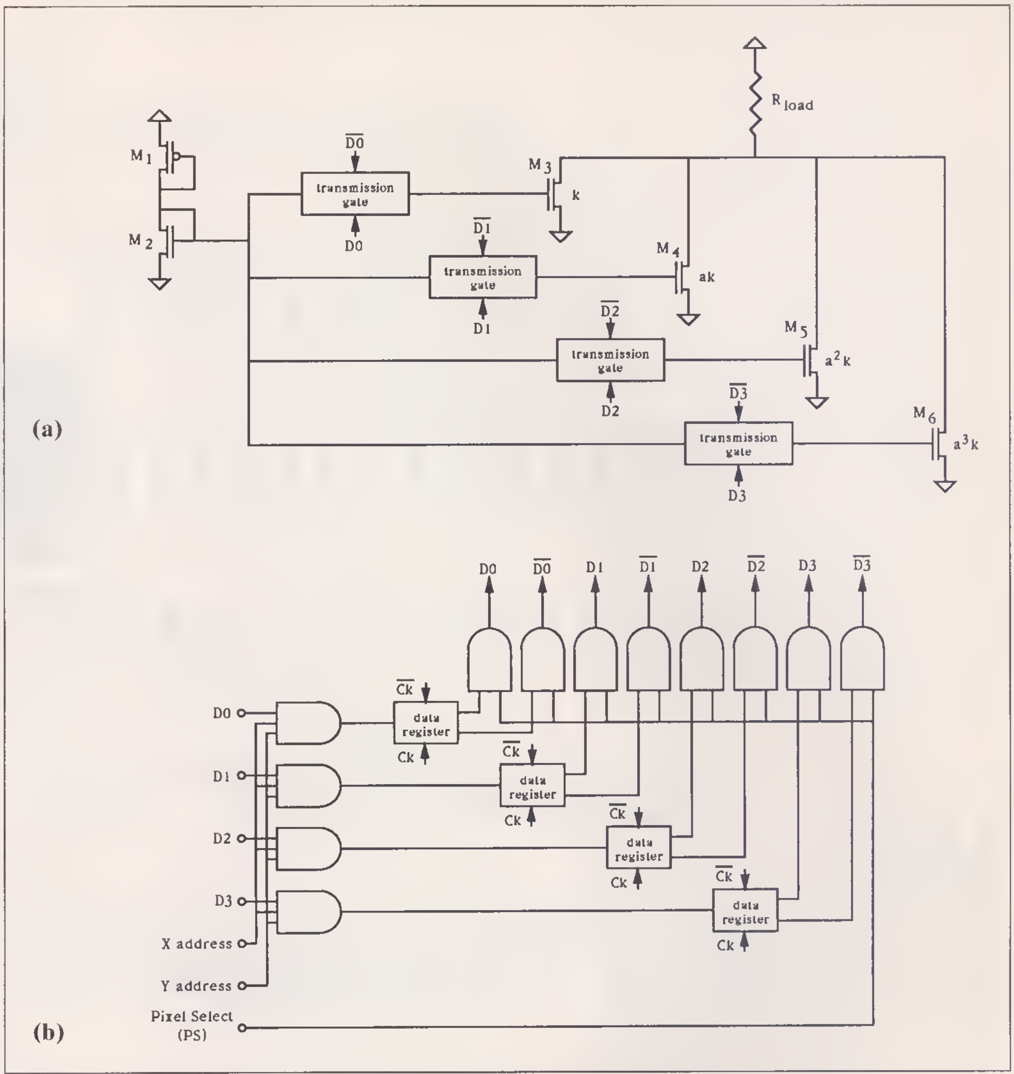

Figure 4. Schematic diagrams of the thermal-pixel driver (a) and control circuit (b) used for each thermal pixel in the wafer-scale thermal-pixel array.

Most of the circuit area is consumed by the array of thermal pixels interleaved by buses linking all system components. The array is much larger than the size of the target display. As Figure 3 shows, each thermal pixel is composed of a 2 × 2 array of subpixels, two identical blocks of pixel-control circuitry, and a set of buses. A working pixel requires two of the subpixels and one control circuitry block. Thus, overall local redundancy is 100 percent. All the subpixels lie on equally spaced grid points, although the size of the *x* and *y* separations need not be the same.

In the cell-harvesting process, we test all thermal pixels (currently using probe pads) and identify the proper area for the highest yield subarray of thermal pixels. We fix the exact row and column subpixel pairs needed for a given pixel address by the choice of the physically used pixel start point (see Figure 2). A more global sparing of pixel-control circuitry between neighboring pixels is feasible because pixel-control blocks lie between rows of thermal pixels. We can laser-link a control block to either the top or bottom row, and right or left column of thermal pixels. Figure 3 also shows the local redundancy of V_{dd}, V_{ss}, clock, data, pixel select, and address buses.

Figure 4 schematically diagrams the circuitry associated with each subpixel. The pixel driver (Figure 4a) uses six metal-oxide semiconductor field-effect transistors (M_1 through M_6) in a current mirror configuration and four transmission gates that control the output current. The currents are controlled by transistor size, which is calculated using Spice models. For linear intensity range, exact thermal modeling of the subpixel heat flow determines the 16 current levels.

The control circuit associated with each pixel (Figure 4b) includes four three-input NAND gates (D0 through D3) that latch the intensity word into a data register for storage. From here, the data controls the drivers' transmission gates when pixel select is activated. Data can be written to the pixels while they are activated for display. Each working subpixel draws relatively high current (2 milliamperes maximum), so only a subset of elements can emit at once. The variable duration of the pixel select signal from the pixel scan generator enables different scanning schemes for dynamic scene simulation.

Our current projection is for a full system that can have about 250 pixels on simultaneously. At this level of operation, the system would draw about 0.5 ampere, and the wafer would stay cool enough to avoid interference with the pixel emissions. The scan generator controls the physical origin point of the displayed pixels.

Figure 3 shows the design flexibility in local sparing of transducers and driver circuits. Any combination of two subpixels constitutes one composite thermal pixel, so many forms of defect avoidance are possible. Because we choose two subpixels from the cluster of four to make a composite pixel, we can easily alleviate the problem of small-area local defects. The redundancy in control circuitry also lets us laser-link control blocks from one pixel site to subpixels at another site. Such links help us circumvent larger defects that may wipe out two neighboring transducers or control blocks, while we maintain proper transducer position in the array. The x, y address select lines programmed for the pixel-control circuitry are those of the physical address of the pixel itself, not of the control section.

We have fabricated small arrays of thermal pixels and control circuitry in 3-micron CMOS technology and measured them. We are now fabricating revised versions and are testing new pixel designs that reduce power consumption and increase emission uniformity. Larger pixel arrays and laser-linked test chips have been sent for fabrication, and we have planned experiments on the redundancy scheme. We are performing etches of full wafers of suspended plate structures to address the wafer-scale yield question.

Like many wafer-scale projects, our research is still in the early stages, but we have fabricated the subelements of at least one design. The results of our experiments will determine the feasibility of a full WSTA implementation. Other WSTAs we are studying include a temperature-sensor array and a smart vision sensor. ■

References

1. *Proc. Int'l Conf. Wafer-Scale Integration*, IEEE CS Press, Los Alamitos, Calif., Order No. 2126, 1991.
2. D.M. Stewart, "The Economics of Laser Repairing the One Megabit Dynamic RAM," *Technical Proc. Semicon/West*, 1984.
3. A.H. Anderson, J.I. Raffel, and P.W. Wyatt, "Wafer-Scale Integration Using Restructurable VLSI," *Computer*, Vol. 25, No. 4, Apr. 1992, pp. 41-47.
4. M. Parameswaran et al., "A New Approach for the Fabrication of Micromechanical Structures," *Sensors and Actuators*, Vol. 19, 1989, pp. 289-307.
5. G.H. Chapman, J.M. Canter, and S.S. Cohen, "The Technology of Laser Formed Interactions for Wafer-Scale Integration," *Proc. Int'l Conf. Wafer-Scale Integration*, IEEE CS Press, Los Alamitos, Calif., Order No. 1901 (microfiche only), 1989, pp. 21-29.
6. C.M. Meijer, "A Low-Power Easy to Calibrate Temperature Transducer," *IEEE J. Solid-State Circuits*, Vol. SC-17, 1982, pp. 609-613.
7. D.R. Briglio, *Characterization of CMOS Magnetic Field Sensors*, master's thesis, Univ. of Alberta, Edmonton, Canada, 1984.
8. C. Mead, *Analog VLSI and Neural Systems*, Addison-Wesley, Reading, Mass., 1989.
9. K.E. Petersen, "Silicon as a Mechanical Material," *Proc. IEEE*, Vol. 70, 1982, pp. 420-457.
10. M. Parameswaran et al., "Micromachined Thermal Radiation Emitter from a Commercial CMOS Process," *IEEE Electron Devices Letters*, Vol. EDL-12, No. 2, 1991, pp. 57-59.
11. L. Burriesci and R. Keezer, "A Dynamic RAM Imaging Display Technology Utilizing Silicon Blackbody Emitters," *Proc. SPIE*, Vol. 765, 1987, pp. 112-122.

Glenn H. Chapman is an associate professor in the School of Engineering Science at Simon Fraser University. From 1980 to 1990 he was a member of the research staff at MIT Lincoln Laboratory, working as the principal researcher in the development of laser redundancy techniques for the wafer-scale laser-restructurable VLSI project. He holds three patents on laser integrated-circuit redundancy technology, and three on ion-beam modification of dielectric films into conductors.

Chapman received his BSc in engineering physics and his MSc in engineering physics from Queen's University, Kingston, Ontario, in 1972 and 1975, respectively. He received his PhD in physics from McMaster University, Hamilton, Ontario, in 1982. He is a senior fellow of the British Columbia Advanced System Institute, a member of the IEEE and the IEEE Computer Society, and chairman of the Electronics Group at Simon Fraser University.

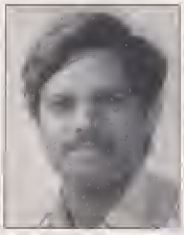

M. Parameswaran is on the faculty of the School of Engineering Science at Simon Fraser University, where he established a micromachining facility to pursue research in micromachined sensors and actuators. He is currently involved in micromachining and sensor research for applications such as integrated infrared emitters, integrated accelerometers, flow sensors, and resonant biomass sensors.

Parameswaran received his BE with honors in electronics and communication engineering from the University of Madras, India, in 1982; his MSc in electrical engineering from the University of Alberta, Edmonton, Canada, in 1985; and his PhD in electrical engineering from the University of Alberta in 1990. He is a member of the IEEE.

Marek Syrzycki is an associate professor in the School of Engineering Science at Simon Fraser University. His research interests are integrated-circuit design, integrated transducers, novel semiconductor devices and circuits, sensor fusion, and computer-aided design.

Syrzycki received his MSc in electronic engineering from the Warsaw University of Technology in 1970 and his PhD in electrical engineering from the Institute of Electron Technology, Research and Production Centre of Semiconductor Devices, Warsaw, in 1978.

Readers can contact Glenn Chapman at Simon Fraser University, School of Engineering Science, Burnaby, B.C., Canada V5A 1S6, e-mail glennc@cs.sfu.ca.

Carnegie Mellon University
Software Engineering Institute

We're improving the practice of software engineering.

Located at Carnegie Mellon University in Pittsburgh, Pennsylvania, the Software Engineering Institute provides leadership in advancing the state of the practice of software engineering.

Motivated professionals with experience in software process, software risk management, technology transition, program management, education, and network security are essential to our mission in 1992.

To get information or to learn about our subscription program, educational courses, or annual symposium call (412) 268–5800.

Care to join us?

The Software Engineering Institute is an affirmative action, equal opportunity employer sponsored by the U.S. Department of Defense.

Leadership in Software Engineering

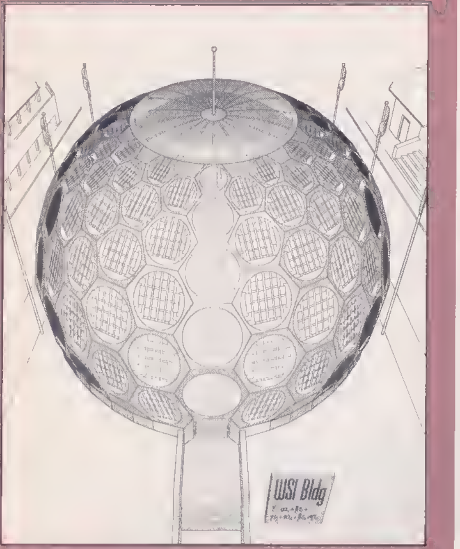

Prospects for WSI: A Manufacturing Perspective

Wojciech Maly, Carnegie Mellon University

For 20 years, the drive toward smaller feature size has dominated the evolution of microelectronics. The cost analysis presented here suggests a change in this trend and better prospects for WSI.

The concept of wafer-scale integration is already more than 20 years old. Yet despite many clever innovations, years of development, and millions of dollars invested in WSI, a large digital system implemented on a single silicon wafer is still more an academician's dream or a costly exhibit on the wall of a government-funded research lab than a commercially viable product delivering all the theoretically possible performance, reliability, and economic gains.

Why this discrepancy between hopes and realities? When will WSI finally prove its usefulness? Do recent promising announcements about WSI memories indicate the beginning of a new era of WSI microelectronics?

These and many similar questions have been asked often by academic visionaries, industrial strategists, designers, and others who want to understand the future of the still rapidly evolving microelectronics field. Investigations undertaken to answer these questions have usually focused on performance and yield. In this article, the focus is on manufacturing cost — an aspect of WSI that has been more or less ignored in the past, especially in research-oriented circles. But with the economic realities of our time, manufacturing cost may prove the most important factor in determining the future of WSI. This article derives a manufacturing cost model to analyze the past and predict the future of WSI.

Past progress in microelectronics

A number of vividly obvious trends characterize progress in the past 20 years of microelectronics. One is a constant decrease in the minimum feature size. Corresponding with this trend has been a constant increase in, for instance, the memory that can be implemented on a single die.[1,2] However, while the increase in memory is more or less exponential, the decrease in feature size tends to be linear. The historical data in Figure 1 show a tendency for feature size to decrease by 0.2 μm every two years.

Will these trends continue in the future? Can feature size decrease along the straight line shown in Figure 1 to the year 2000? The answer is no, for a variety of reasons. The difficulties in optical lithography below 0.3 μm, the need to scale

0018-9162/92/0400-0058$03.00 © 1992 IEEE

down voltage supplies (which decreases the signal-to-noise ratio), and the tunneling and breakdown effects in a very thin oxide are among the fundamental physical obstacles that have been discussed in the literature in detail.[3] There is, however, an additional obstacle: the economic infeasibility of fabricating ever more complex circuits with ever more expensive processes.[4] In just a few years, this may become a dominant obstacle to achieving smaller feature sizes.

Figures 2-4 explain the extent and nature of this problem by showing increases in manufacturing costs as functions of minimum feature size, denoted by λ.* Figure 2 depicts the cost of lithography as a function of λ. Note that this curve has an exponential nature. The relative cost of lithography shown in Figure 2 can be approximated by a function $X^{\alpha(\lambda)}$, where the parameter X represents a *rate of cost increase* and $\alpha(\lambda)$ is a simple function of λ. For example, in Figure 2, $X = 1.8$ means that the cost from one generation of technology to the next increases 180 percent. (The decrease in feature size between generations is assumed to be 0.2 μm.)

With a slightly different value for parameter X, the same function can model the cost of the fabrication line and the cost of the wafer, as shown in Figure 3.

In general, studies of the available data describing the cost of fabrication lines, the pieces of processing equipment, and the resulting cost of the manufactured wafer indicate that X has values in a range between 1.2 and 2.5. Later, this article will show that the value of X is extremely important in explaining microelectronic trends, past and future.

There are many reasons for the exponential increase in manufacturing costs: changing market needs, growing research and development costs, time required for equipment development, international competition, and technological complexity. These are among the main factors that force investors to obtain increasing amounts of money to

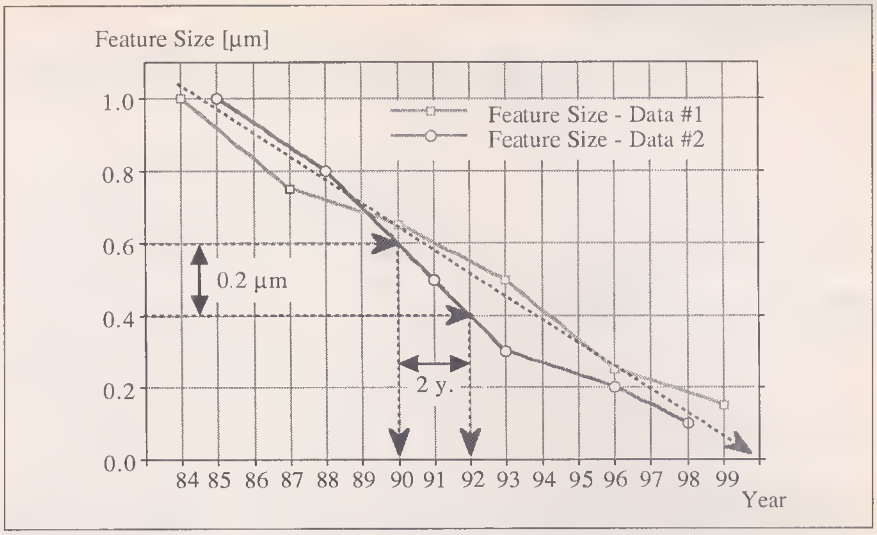

Figure 1. Minimum feature size as a function of time.[1,2]

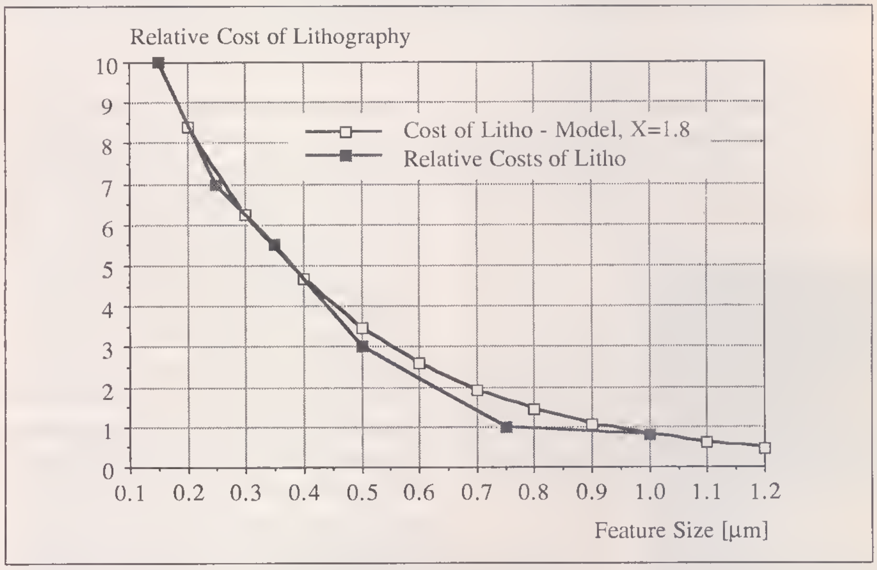

Figure 2. Relative cost of the lithography as a function of the minimum feature size.[1]

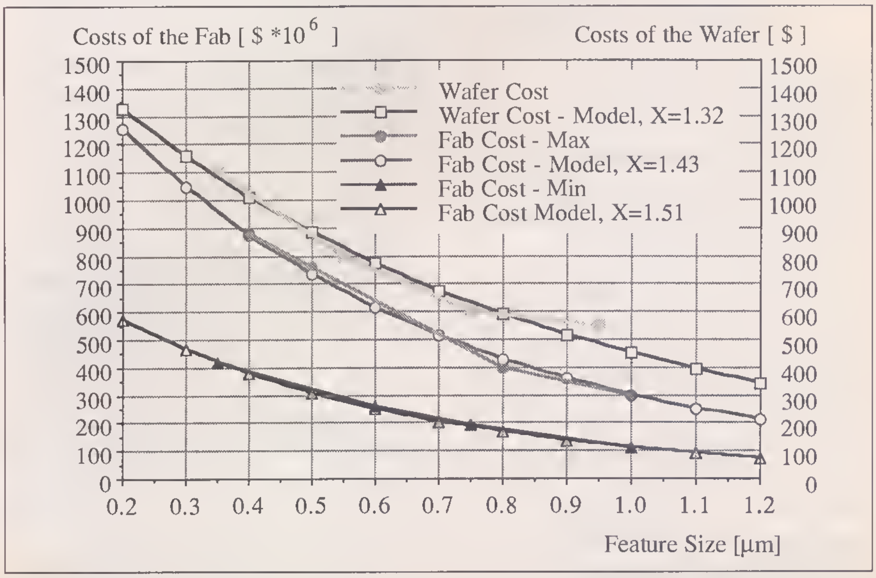

Figure 3. Cost of fabrication line and cost of the 6-inch manufacturing wafer as a function of the minimum feature size.[1] **Data is normalized to represent a fabrication line with 10,000 starting wafers per month.**

---

*In an IC, polysilicon lines usually have minimum width. Since these lines form gates of the MOS transistors, the line width is equivalent to the transistor's channel length. Consequently, minimum feature size, λ, is equal to the length of the shortest MOS transistor channel.

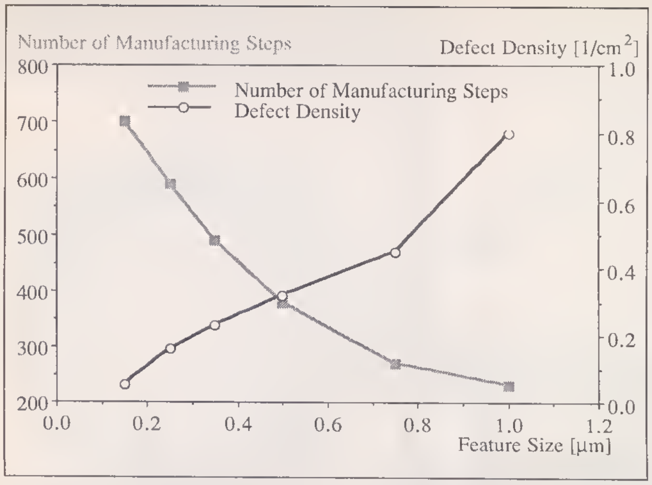

Figure 4. Complexity of the process (expressed in terms of the number of manufacturing steps) and quality of the process (expressed in terms of the number of defects per square centimeter) as functions of minimum feature size.[1,6]

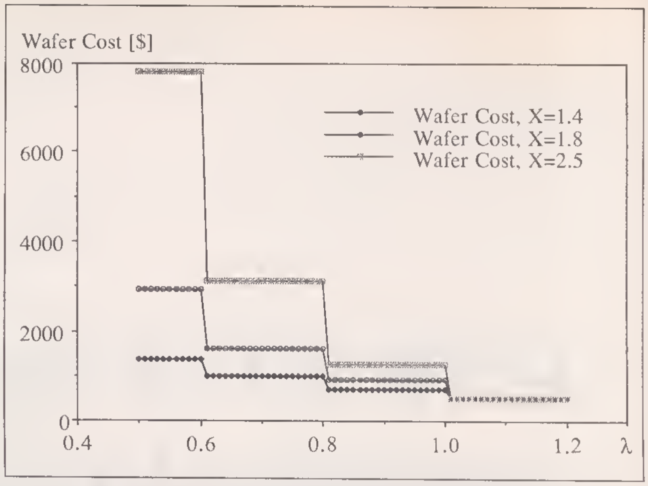

Figure 5. Cost of the 6-inch manufacturing wafer computed for X = 1.4, 1.8, and 2.5, assuming that such a wafer costs $400 if fabricated with 1-μm technology. It is also assumed that the cost of the wafer is not a function of the die size.

meet growing process costs. Figure 4 explains these cost trends by showing the number of manufacturing steps and the defect density as functions of λ in the fabrication of memories, according to predictions published in the literature.[2] Notice that the increase in the number of manufacturing steps is faster than linear. Even more important, however, is the rate of the required decrease in defect density.[5] The problem is that defect density is a strong function of the defect size,[3] which can be approximated with the following formula:[6]

$$f(R_d)=\begin{cases}\dfrac{2p\,R_d}{(p+2)\,x_0^2} & \text{for } 0\le R_d\le x_0\\[2ex] \dfrac{2p\,x_0^p}{(p+2)\,R_d^{p+1}} & \text{for } x_0\le R_d\end{cases}\qquad(1)$$

where $f(R_d)$ is the defect density, R_d is the defect radius, and p is a parameter with values in the range between 2 and 4.

To explain the difficulty caused by the phenomenon characterized by Equation 1, let us assume that a defect is a region (disk) of extra or missing material occurring in one of many layers of an integrated circuit. For instance, if such a "disk" is a spot of a conductor deposited in a conducting layer, then it can short out some nonequipotential regions in this layer, causing IC malfunction. One can see that any defect with a radius larger than half the distance between the edges of two electrically distinct regions can trigger an electrical malfunction of the circuit.

From this, it is evident that the decrease of the minimum feature size increases the fraction of the defects that can cause a short or a break in the circuit. Consequently, the decrease of the minimum feature size may lead to a drastic decrease of yield. The strength of this effect is related to the value p in Equation 1: the larger the value of p, the more rapid the increase in the chance of yield loss. Therefore, to maintain an acceptable level of yield, a decrease in the minimum feature size must be followed by a rapid decrease in the defect density — which is difficult because of more frequently occurring small defects. In other words, an adequate contamination-prevention strategy must be very aggressive to eliminate a rapidly growing number of smaller defects. Of course, the more aggressive the contamination control must be, the more expensive it is.

Further, contamination control becomes very difficult when defect sizes are in the range approaching the limits of optical-based contamination-monitoring techniques. Smaller devices are also sensitive to many kinds of contamination that cannot be detected unless very sophisticated analytical methods are applied. Hence, in general, the increasing complexity of the process and the necessary cost of contamination control must contribute to the rapid growth of manufacturing costs. In the future, one can therefore expect more rapid increases in the manufacturing costs and consequently larger values of parameter X in the wafer cost model.

Cost model

Growth in the cost of manufacturing is an inherent phenomenon. But is it really a negative factor limiting the progress of microelectronics? The answer to this question is not simple, because higher manufacturing costs can be offset by rapidly growing functionality of the fabricated ICs. This section examines this problem in more detail by analyzing the relationship between manufacturing costs and development patterns known from the past. The analysis uses a simple cost model that expresses manufacturing cost per transistor in terms of such variables as minimum feature size, die area, previously introduced parameters X and p, and a few others.

The basis for this modeling effort is the wafer cost model discussed in the previous section. The model development assumed that wafer cost is a function of the minimum feature size, λ, normalized with respect to the cost of a 6-inch wafer fabricated with feature size λ = 1 μm. Wafer cost was expressed in terms of the rate of the cost increase, X, introduced in the previous section. Note that any new generation of technology requires investment in new, better equipment and in a necessary infrastructure. Since the cost of technology development and the cost of new equipment constitute a large portion of the manufacturing costs, the cost per wafer must be described using a step function — with the step increase of the cost occurring only for λ-changing borderlines between generations of technologies. It was decided to locate borderlines between technologies considered in this article at minimum feature sizes equal to 1.2, 1.0, 0.8, 0.6, and 0.4 μm.

Using these assumptions and available cost data discussed in the previous section, a function $C_{\text{waf}}()$ was built. This function was used, for instance, to compute the cost of a 6-inch manufacturing wafer for three X values. Figure 5 shows the results. Note that X is used as a parameter but with values close to the values extracted from the real data presented in Figure 3. Therefore, in subsequent sections of this article, the wafer cost model, $C_{\text{waf}}()$, is assumed to depict a realistic range of possible wafer cost options.

Wafer size, minimum feature size, and cost of the wafer can be used to determine the maximum number of transistors that can be fabricated on a single wafer, and consequently the cost of manufacturing per fabricated transistor. But to figure out potential manufacturing profits requires knowing the cost per transistor on an operational (or simply defect-free) chip — not the cost per potentially working transistor. Hence, the figure of merit to discuss here is a ratio of wafer cost C_{waf} divided by a product $N_{\text{tr}} \times N_{\text{chip}} \times Y_{\text{par}} \times Y_{\text{fun}}$, where N_{tr} is the number of transistors on a chip, N_{chip} is the number of chips per wafer, and Y_{par} and Y_{fun} are parametric and functional yields, respectively. In other words, the needed figure of merit should have the following form:

$$\text{Cost}_{\text{tr}}(X,\lambda,C_0,D_0,p,x,R_j,\sigma,nl) = \frac{C_{\text{waf}}(C_0,X,\lambda)}{N_{\text{tr}}\,N_{\text{chip}}(R_j)\,Y_{\text{par}}(\sigma,\lambda,nl)\,Y_{\text{fun}}(D_0,\lambda,p)} \qquad (2)$$

where

X, as previously described, is an investment rate factor;

λ is a minimum feature size;

C_0 is the cost of a 6-inch wafer fabricated with 1-μm technology;

D_0 is a defect density;

p is a parameter of defect size distribution;

x is the length of an edge of the square die;

d is the distance forming the base for all geometrical design rules;

R_j is the radius of the wafer;

σ is the tolerance of the total overlay error possible with the given lithography; and

nl is the number of masks in a given process.

Full explanation of all components used to build the formula in Equation 2 exceeds the scope of this article. Note, however, that its basic element, function $C_{\text{waf}}(C_0, X, \lambda)$, is like the function used to produce the curves shown in Figure 5. It is assumed, therefore, that $C_{\text{waf}}(C_0, X, \lambda)$ models the realities with the accuracy discussed at the beginning of this section and the previous one. In other words, one can assume that C_{waf} can well describe the cost increase if the value of X is known. Functions $N_{\text{chip}}(R_j)$, $Y_{\text{par}}(\sigma, \lambda, nl)$, and $Y_{\text{fun}}(D_0, \lambda, p)$ were derived using simple geometrical relationships[4] and yield models.[3] Consequently, the computation of parametric yield loss in Equation 2 assumes that the actual distance between the two closest edges in the IC must be greater than zero if their nominal distance is d and if the "misplacement" of the actual edges is Gaussian. It is assumed that such a misplacement is characterized by σ, which can also be interpreted as a standard deviation of the total overlay error.

The functional yield loss is represented by a function derived from the Poisson model. This function models a defect-sensitive area (called a critical area) in terms of defect size, discussed earlier. (Consequently, Y_{fun} is related to the defect size distribution parameter p.) Hence, despite its simple form, Equation 2 models cost per transistor with sufficient accuracy to depict the relationships between its parameters in a realistic manner.

The following sections use Equation 2 to study basic design trade-offs and to predict future trends in microelectronics in terms of a few key parameters: X, p, and σ.

Past development trends in microelectronics

Several market-driven objectives and limitations have shaped the evolution of microelectronics. The most important have been

(1) a need to introduce a new product to the marketplace at the earliest possible moment,

(2) a need to provide more function-

Carnegie Mellon University
Software Engineering Institute

Career Opportunities

Make a difference.

We are looking for senior staff who have–

5 years of successful project management experience

5 years of proven track record of involvement in software development or software acquisition

5 years of management or development of large-scale, embedded, real-time systems

Strong communication and leadership skills

Master's or Ph.D. degree

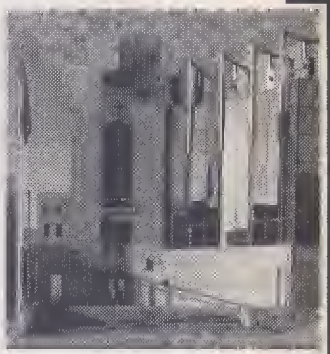

We have openings for qualified applicants in the following areas related to software–

Process assessment

Process maturity modeling

Process definition

Risk management

Software capability evaluation

Send resumé to

Sally Miller
Software Engineering Institute
Carnegie Mellon University
Suite 114
Pittsburgh, PA 15213–3890

The Software Engineering Institute is an affirmative action, equal opportunity employer sponsored by the U.S. Department of Defense.

ality in any new generation of electronic systems entering the marketplace,

(3) a need to decrease the price of any new generation of electronic system entering the marketplace (unless it can provide drastically new utilities or breakthrough performance), and

(4) a tendency to minimize cost per fabricated product by increasing volume of fabrication.

Of course, all four of these objectives can be derived from one underlying need: to maximize profit margins.

How do these needs and tendencies correlate to the known patterns of evolution presented in this article? To answer this question it is first necessary to note that items (2) and (3) can be translated into a requirement for a constant increase in the number of transistors on a single chip and, possibly, a decrease in the cost of fabrication per transistor for each new generation of semiconductor technology. Progress of this nature requires a decrease in the minimum feature size, an increase in the die area, or any combination of the two.

To better characterize trade-offs between die size, minimum feature size, and cost of fabrication, one can use the cost model summarized in the previous section and apply it to generate diagrams like the one shown in Figure 6. In this diagram, straight lines in a two-dimensional space, $\lambda \times x$, represent IC dies having the same number of transistors. (Note that since a line in the $\lambda \times x$ plane is represented by the equation $\lambda = ax$, then the λ / x remains constant for all points on the line.) Hence, the dashed straight line shown in Figure 6 indicates changes in the minimum feature size and die size that preserve the number of transistors on the die. The solid lines represent constant cost per transistor; they are computed by solving for λ and x in the equation

$$\mathrm{Cost}_{tr}\,(X, \lambda, C_0, D_0, p, x, R_j, \sigma, nl) = C_i \text{ for } i = 1, 2, \ldots,$$

where C_i is an assumed cost per transistor.

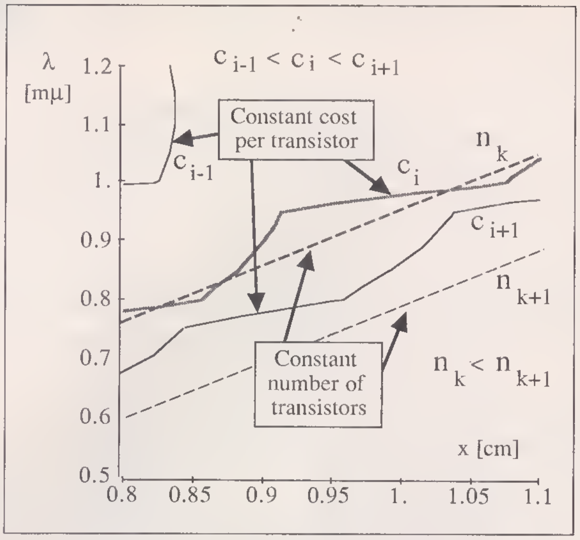

Figure 6. Cost trade-offs between minimum feature size, λ, and edge of the square die, *x*.

Of course, the parameters of the cost model in Equation 2 determine the shape of these piecewise linear constant-cost contours. Notice that an increase in the number of transistors achieved by decreasing the minimum feature size is equivalent to movement in the $\lambda \times x$ plane along any vertical line pointing in the negative direction of the λ axis. An increase in the number of transistors achieved by increasing the die size is equivalent to movement along any horizontal straight line in the positive direction of x.

Let us consider now the set of contours presented in Figure 7. This set was computed for parameters known in the past to depict typical situations in the semiconductor industry and representing an optimistic view of the future: $X = 1.3$ indicates a 30-percent increase in the cost per generation, $p = 2.0$ represents the defect size distribution achieved in the past by using a traditional contamination-prevention strategy, and $\sigma = 0.2$ represents the overlay tolerance achieved with the existing lithography technique. Notice that the orientation of the constant-cost contours is almost vertical (with higher cost contours at the right-hand side of the drawing). One can see that in such situations increasing the number of transistors by decreasing the minimum feature size should not increase the cost. However, increasing the die size *will* increase the cost. Hence, data obtained from the model suggests that in the past the semiconductor industry had no choice but to decrease the minimum feature size.

Does this conclusion agree with the actual trends observed in the past? Yes. The evolution of microelectronics manufacturing has shown, first of all, a constant or even decreasing cost in manufacturing per bit of memory.[7,8] This was achieved by the decrease in λ and enhanced by the application of redundancy. Such a tendency has provided evidence that minimization of feature size is a viable solution not only for memories but also for microprocessors and eventually for ASIC chips.

Secondly, since equipment development had to be oriented toward high-volume manufacturing, it became dominated by specific needs dictated by the memory segment of the semiconductor industry. These two trends combined have generated favorable conditions for the development of

- high-resolution (rather than large-field) lithography,
- processes oriented to multimask and multistep small transistors (rather than fast interconnects), and
- cleaner (rather than cheaper) manufacturing operations.

They have also determined momentum in the rate of cost increase — that is, in the value of *X* discussed in this article. Consequently, the entire semiconductor industry has become oriented toward a continuous decrease in minimum feature size, while die size has increased very slowly.

This picture seems consistent with the actual evolution and status of the WSI field. To better see the complexity of this evolution, note that the cost model (Equation 2), with the parameters used to produce Figure 7, explained trends in equipment and process development, but cannot indicate the feasibility of WSI. The problem is that the cost model does not describe cost per transistor in a WSI system, which can benefit very strongly from, for instance, redundancy techniques.

Nor does the model account for the expense of implementing WSI systems with the lithography developed for very large scale integration (VLSI) and ultra

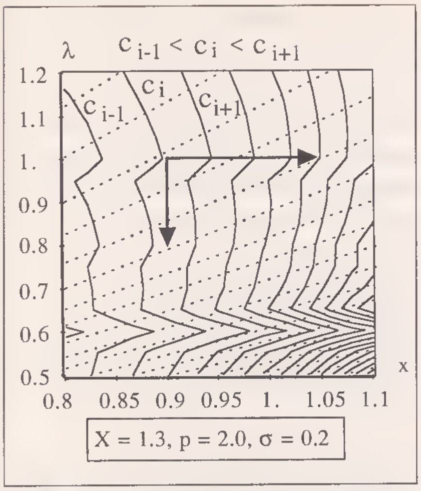

Figure 7. Set of cost contours computed from the cost model (Equation 2) for parameters characterizing semiconductor technologies in the late 1980s.

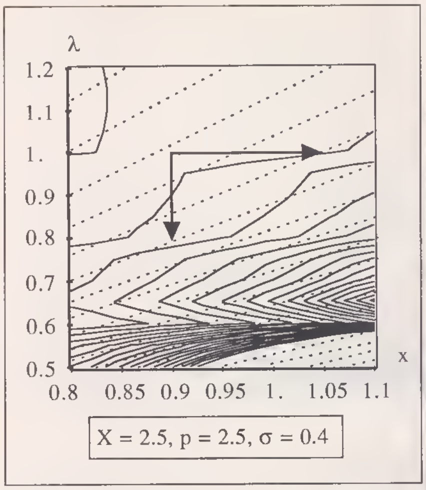

Figure 8. Set of cost contours computed from the model (Equation 2) for a set of parameters characterizing the "worst case" scenario.

large scale integration (ULSI). The increase in wafer cost unaccounted for in Equation 2 is caused by the fact that WSI can either use steppers with expensive reticles, or a very inefficient (and therefore also expensive) direct-write E-beam lithography. If the reticles are used, two options are available: either a single set of masks, which seriously limits the system architecture, or a number of different sets, which dramatically increases the wafer cost* (unless they are produced in large volume).

Hence, a design decision to use a single wafer substrate — assuming that long-range progress must be driven by market forces and by equipment and process development momentum indicated in Equation 2 — would not make sense. The only exception would be if WSI provided performance impossible to achieve in any other way. Unfortunately, due to inherent limitations of the standard interconnect technology (developed for small dies), such gains are unlikely. Only special architectures would benefit from WSI characteristics without suffering from interconnect- and yield-related shortcomings.

The history of WSI confirms these conclusions. WSI did not become a mainstream technology and had to remain a "potential option" that reused technologies, design tools, and equipment developed for VLSI and ULSI.

Future evolution

The cost model (Equation 2) has proven useful in explaining the evolution of microelectronics thus far. This section uses it to predict future trends both in microelectronics as a whole and in WSI specifically.

To make such predictions, however, it is first necessary to discuss again the key parameters of the wafer cost model. As mentioned earlier, the parameters used to compute the constant-cost curves shown in Figure 7 represented past trends and an optimistic view of the future. But the survey of past progress in microelectronics gives little basis for such optimism. The cost increase rate X at the beginning of the current decade is known to be higher than 1.3, and it may go well above 2 when new lithography and contamination-control strategies are implemented. The defect size distribution can worsen, with p shifting closer to 3. The number of processing steps might also increase drastically. Lithography quality (represented in the model by σ) can become a bigger problem than has been anticipated. Hence, to predict the future, one should rely not only on the "best case" scenario, represented in Figure 7, but also on the possible "worst case" scenario, such as the one depicted in Figure 8.

---

*Notice that the lack of an adequate (large-field) lithography technique can offset all benefits from the application of redundancy.

Career Opportunities

Carnegie Mellon University
Software Engineering Institute

Make a move.

We are looking for product planners who have–

10 years of successful experience with software-intensive projects

Product planning or marketing experience

Government or industry client contact experience

Working knowledge of software development

Strong communication skills

Master's degree (preferred)

We have openings for qualified applicants in the following areas–

Product and services

Organizing & planning

Software process

Engineering management

Send resumé to

Sally Miller
Software Engineering Institute
Carnegie Mellon University
Suite 214
Pittsburgh, PA 15213–3890

The Software Engineering Institute is an affirmative action, equal opportunity employer sponsored by the U.S. Department of Defense.

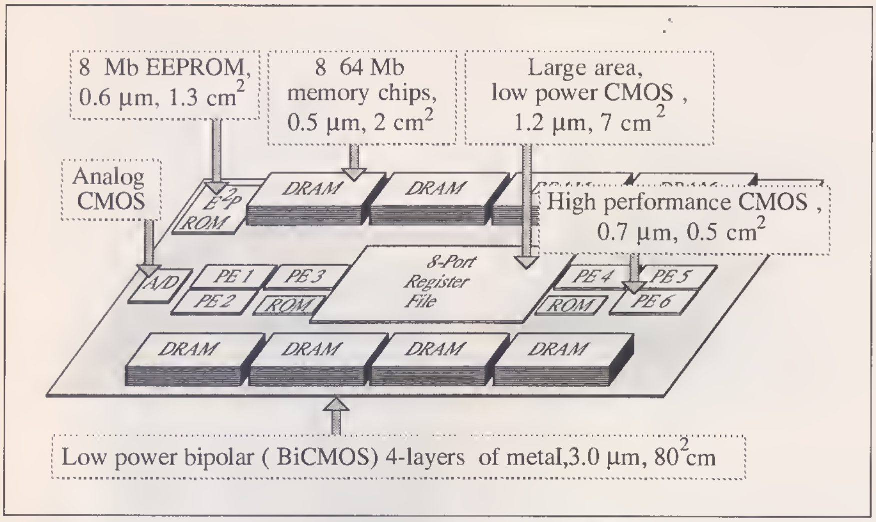
8 Mb EEPROM, 0.6 μm, 1.3 cm²
8 64 Mb memory chips, 0.5 μm, 2 cm²
Large area, low power CMOS , 1.2 μm, 7 cm²
Analog CMOS
E²P ROM
DRAM
DRAM
High performance CMOS , 0.7 μm, 0.5 cm²
A/D
PE 1
PE 2
PE 3
ROM
8-Port Register File
PE 4
PE 5
ROM
PE 6
DRAM
DRAM
DRAM
DRAM
Low power bipolar (BiCMOS) 4-layers of metal,3.0 μm, 80 cm²

Figure 9. Optimal solution of the cost-performance trade-off.

Figure 8 presents a new set of contours, computed for $X = 2.5, p = 2.5$, and $\sigma = 0.4$. Notice that in this case the cost picture is dramatically different from the situation depicted in Figure 7. A decrease in minimum feature size now causes an increase in the cost per transistor. More importantly, an increase in die size may produce a smaller cost increase than an equivalent decrease in minimum feature size. This indicates that the evolution pattern successful in the past may no longer be valid and that one should consider large-area dies rather than ULSI to increase the number of transistors placed on a single die!

Late Magazines?
No Magazines?
Membership Status Problems?
No Answers To Your Complaints?
Let your Computer Society Ombudsman cut through the red tape for you.

Ombudsman
IEEE Computer Society
10662 Los Vaqueros Circle
PO Box 3014
Los Alamitos, CA
90720-1264

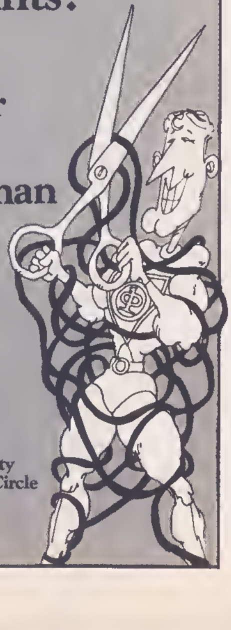

Notice that if this "worst case" scenario is realistic, then several possible changes could affect the entire semiconductor industry. First, the ability to maintain the constant cost per bit of memory will vanish, slowing the trend towards a smaller feature size. In addition, segments of the industry fabricating silicon with a density lower than the density of memories (and without redundancy) might see no benefit in going to the ULSI regions at all. Consequently, the volume of the silicon fabricated with the most advanced technologies (smallest λ) would decrease. In turn, the equipment industry might find it beneficial to pay more attention to processes and technologies oriented towards larger dies. Ultimately, larger dies could become even more economical by getting better support from the infrastructure than in the past.

It must be stressed here that such conclusions are derived directly from a simple but realistic cost model. In addition, since the parameters used in computations of the data in Figure 8 are not unlikely, these predictions should be considered quite realistic.

An alternative solution

The information presented so far indicates that neither ULSI nor WSI currently offers ways to decrease the per-transistor costs of fabricated devices. What are the other options? The answers are simple:

(1) develop a more cost-effective way to fabricate ULSI wafers,

(2) develop ways of using expensive fabrication technologies more effectively, or

(3) direct the equipment industry toward the needs of large-area devices as progress in smaller feature size technology slows down.

Option (1) can be effective for high-volume memory-like products but not for lower density and redundancy-free applications. Hence, after full development of 200-mm wafer processes, industry will pay more attention to option (2).

Option (2) is probably the most promising and the most likely to materialize soon. Full analysis of this option exceeds the scope of this article, but Figure 9 illustrates it well by depicting a concept of 2.5D integration scheme derived from ideas postulated by McDonald[9] and Tewksbury.[10] This concept can be viewed as an advanced version of the recent multichip module (MCM) packaging strategy.

It is envisioned that this hypothetical system will respond directly to the cost limitations of VLSI technologies discussed in this article. Notice its basic characteristics. The system could be assembled on a large-area active substrate. The technology of such a substrate could be optimized for yield, power, and speed of the interconnect. This substrate could dissipate a large percentage of the total power and could be cost-effective if fabricated with relaxed design rules in stepper-free, interconnect-oriented technology. The performance-critical system components could be fabricated separately on fabrication lines oriented toward high volume and high performance. Then they could be attached to the active substrate with the rapidly maturing flip-chip technology. This way only those system elements that really require ULSI technology (for example, data path) would be fabricated with the most expensive technologies.

At present, the 2.5-dimensional integration, which could take advantage of the flexibility of the proposed hybrid assembly technique, seems the best answer to the rapidly growing cost of manufacturing. In the long run, however, option (3) will probably prevail. This is because MCM technology will become less attractive as new equipment allows cheaper production of larger systems on a common substrate. This option will take longer to evolve and become cost-competitive, but eventually it will lead to commercially successful WSI production.

The manufacturing cost model used in this article to explain major trends in the past 20 years of microelectronics describes the relationships among characteristics of modern manufacturing processes, investment costs to achieve these characteristics, and basic IC parameters, including both die size and minimum feature size. Results from this model indicate that it is not possible to continue progress in microelectronics through the minimization of feature size. Therefore, in the longer perspective, it is anticipated that the drive toward larger dies will gain momentum and lead gradually toward WSI.

Data from the analysis presented here also indicates that manufacturing costs will keep WSI from becoming practical in the immediate future. As an alternative, active-substrate flip-chip MCMs may provide both the performance gain and cost efficiency required. ■

Acknowledgments

The author thanks Ron Rohrer, Derek Feltham, Hans Heineken, and Jitendra Khare for their help in preparing this article.

The work was supported in part by National Science Foundation Grant No. MIP-8822805.

References

1. C.R. Deninger, "Fabs of the 2000," *Proc. Techcon 90*, CRC Press, San Jose, Calif., 1990 pp. 276-278.
2. G.E. Flores and B. Kirkpatrik, "Optical Lithography Stalls X Rays," *IEEE Spectrum*, Vol. 28, No. 10, Oct. 1991, pp. 24-27.
3. W. Maly, "Computer-Aided Design for VLSI Circuit Manufacturability," *Proc. IEEE*, Vol. 78, No. 2, Feb. 1990, pp. 356-392.
4. W. Maly, "Modeling of Lithography-Related Yield Losses for CAD of VLSI Circuits," *IEEE Trans. CAD and Systems*, Vol. CAD-4, July 1985, pp. 166-177.
5. C.H. Stapper, R.M. Amstrong, and K. Saji, "Integrated Circuit Yield Statistics," *Proc. IEEE*, Vol. 71, Apr. 1983, pp. 453-470.
6. V. Menon, "Microcontamination, Defect Measurement and Control in ULSI Manufacturing," *Proc. Techcon 90*, CRC Press, San Jose, Calif., 1990, p. 224.
7. T. Masuchara et al., "VLSI Memories: Present Status and Future Prospects," *Institute of Electronics, Information, and Communication Engineers (IEICE) Trans.*, Vol. E-74, No. 1, Jan. 1991.
8. Y. Tarui and T. Tarui, "New DRAM Pricing Trends: The Bi Rule," *IEEE Circuits and Devices Magazine*, Mar. 1991, pp. 44-45.
9. J.F. McDonald et al., "The Trail of Wafer-Scale Integration," *IEEE Spectrum*, Oct. 1984, pp. 32-39.
10. S.K. Tewksbury, *Wafer-Level System Integration: Implementation Issues*, Kluwer Academic Press, Boston, 1989.

Wojciech Maly is a professor of electrical and computer engineering at Carnegie Mellon University and codirector of Sematech Center for Excellence. He also consults for a number of companies in the US and Europe and is a member of the Technical Advisory Board of Cadence. His research interests are computer-aided design, testing, and manufacturing of integrated circuits.

Maly received an MSc in electronic engineering from the Technical University of Warsaw, Poland, in 1970, and a PhD with honors from the Institute of Applied Cybernetics, Polish Academy of Sciences, Warsaw, in 1975. He is a fellow of the IEEE, a member of the IEEE Solid State Circuit Council, an associate editor of *IEEE Transactions on Semiconductor Manufacturing*, a member of the Editorial Board of the *Journal of Electronic Testing*, and consulting editor for CRC Press.

Readers may contact the author at the Electrical and Computer Eng. Dept., Carnegie Mellon Univ., Pittsburgh, PA 15213.

Career Opportunities

Carnegie Mellon University
Software Engineering Institute

Solve a mystery.

We are looking for problem solvers who have–

Technical problem solving skills

Software development experience

In-depth knowledge of Unix

Strong personal interaction and communication skills

Bachelor's degree

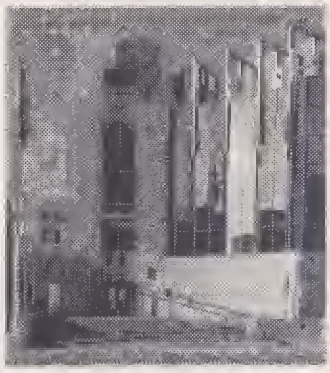

We have openings for qualified applicants in the following areas–

Network security

Computer emergency response

Send resumé to

Sally Miller
Software Engineering Institute
Carnegie Mellon University
Suite 314
Pittsburgh, PA 15213–3890

The Software Engineering Institute is an affirmative action, equal opportunity employer sponsored by the U.S. Department of Defense.

PROJECT OVERVIEWS

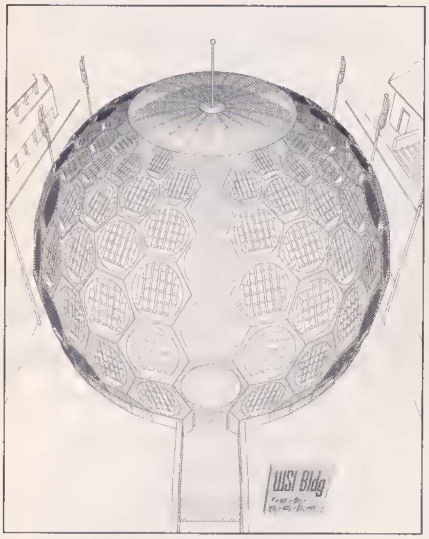

The following overviews report progress in ongoing research in two areas: the use of redundancy in WSI design and the development of a WSI architecture for rapid prototyping.

Wafer-Scale Optimization Using Computational Availability

David L. Landis, Nitin Nigam, and Joseph W. Yoder
University of South Florida, Center for Microelectronics Research,
Tampa, FL 33620-5350

The size of monolithic integrated circuits has generally been limited by photolithographic capabilities and by yield loss associated with defects. Advancing technology has reduced feature size and defect densities, allowing higher levels of on-chip integration. However, the relationship between defects (that is, yield) and cost always limits the maximum economical chip size.

An alternative to nonredundant very large scale integration (VLSI) designs is to develop IC design methodologies and manufacturing technologies that can tolerate defects. This is accomplished through the selective use of redundant components, along with the restructuring of each fabricated device to circumvent its unique defect pattern. This technique is routinely used in RAM design. At the wafer level, successful methods for restructuring designs have been developed and demonstrated at the Massachusetts Institute of Technology's Lincoln Laboratory, at the University of South Florida's Center for Microelectronics Research, and elsewhere.

Given the ability to restructure wafer-level designs, there are different ways to employ redundancy. In this article, we evaluate redundancy by estimating system computational availability over a mission lifetime. We illustrate this technique with two wafer-scale integration (WSI) case studies. The first is a very fine-grained programmable systolic data processor (PSDP) that contains 4- and 8-bit data paths, RAM, and control optimized for signal and data processing applications. By contrast, the Mosaic multicomputer architecture, developed at the California Institute of Technology,[1] represents a less fine-grained homogeneous architecture where each node contains a 16-bit microprocessor and associated RAM and ROM.

Potential benefits of implementing these parallel processing architectures in wafer scale include

- reduced system pin count, which improves system reliability;
- shorter and lower capacitance interconnect, which increases speed and reduces power requirements;
- fewer off-package drivers and smaller interprocessor drivers, resulting in lower power and higher density; and
- fewer parts per system, resulting in a smaller, lighter, and more reliable system.

We begin by considering a simple homogeneous WSI multicomputer design. The limitations of photolithography preclude patterning the entire wafer as a single exposure. Thus, the required step-and-repeat process dictates the organization of a WSI design as a set of repetitive elements called tiles. If the basic WSI tile is a single processor, then the intuitive method is to partition the wafer into as many tiles (processors) as possible. After wafer fabrication, the good processor nodes are connected (harvested) to form a working multicomputer.

When hard restructuring (for example, laser linking) is used to form the connections, then the architecture is fixed. However, hard restructuring is done at the end of production before the system is operational; any component failure after restructuring could lead to complete system failure.

By contrast, when soft (programmable) switches (for example, multiplexers) are used to configure a network of processors, the system can tolerate a runtime failure. In this case, processor-level redundancy can be used for both defect and fault tolerance. Soft switching supports gracefully degradable operation where all fault-free processors are active. Upon failure of any processor, the system can be reconfigured to continue operation with one less processor. Beaudry[2] defines system coverage (c) as the percentage of all possible faults from which a system can recover. In our case, this corresponds to the probability that a WSI system recovers successfully and continues operating in a gracefully degraded mode after a processing element defect or failure.

System performance measures. Interest in dependable computing applications is growing. Accordingly, system engineers have become increasingly sensitive to combined performance and reliability issues and have attempted to achieve optimal performance and reliability through fault tolerance (using detection, diagnosis, repair, and recovery). Two measures of a fault-tolerant system are its reliability and computational availability.

The reliability $R(t)$ of a system is defined as the conditional probability that the system has survived the interval $[0,t]$, given that it was operational at time $t = 0$. This metric reflects the operational lifetime of a system and is inherently a function of the reliabilities of the system's components. When comparing the relative performance of alternative designs, reliability is a limited measure because it does not consider the performance of the resulting system. An alternative used in this study is computational availability, $CA(t)$, which reflects the level of computing performance available at a particular time. It is defined as the expected value of a system's computation capacity at time t:[2]

$$CA(t) = \sum_{i=F}^{N} K * i * P_i(t) \qquad (1)$$

where i represents the system's state (number of working processors), K is the computation capacity of an individual processor, N is the total number of working processors at time $t = 0$, F is the minimum number of working processors, and $P_i(t)$ is the probability that i processors are operational at time t. Beaudry[2] obtains the expression for $P_i(t)$ by solving a continuous time Markov model. The result is:

$$P_i(t) = c^{N-i} e^{-i\lambda t} \binom{N}{i} (1 - e^{-\lambda t})^{N-i} \qquad (2)$$

where c is the coverage, and λ is the processor failure rate.

Redundancy. Selecting the most appropriate form of redundancy is not trivial. This article evaluates two case studies by applying four possible approaches to tile-level redundancy. Both cases use hard restructuring inside a tile to create good tiles, while soft reconfiguration is performed at the global/wafer level to create a good network of tiles.

The structure of the different tiles is summarized as follows:

- each WSI tile contains a single processor and its associated intercell interconnect (1/1);
- each tile is composed of two processors, of which only one can be used (2/1);
- each tile has four processors, of which only two can be used (4/2); and
- each tile again contains four processors, but only three of them can be employed (4/3).

The spare processors in a tile are used for yield enhancement only.

Yield. The probability of x defects occurring within a processing element is calculated using the negative binomial defect distribution model:

$$Pr(Xi = x) = \frac{\Gamma[x+\alpha]}{x!\,\Gamma[\alpha]} \; \frac{\left(\frac{Ad}{\alpha}\right)^{x}}{\left(1 + \frac{Ad}{\alpha}\right)^{\alpha + x}} \qquad (3)$$

where A is the cell area, d is the average defect density, and $\alpha = d^2/\mathrm{var}(d)$ is the clustering parameter.

The predicted tile yield is calculated using a binomial distribution of defective processors within a cell, where the yield of each tile is calculated using the negative binomial model:

$$Y = \sum_{j=0}^{sp} \binom{N}{j} Pr^{N-j} (1 - Pr)^{j} \qquad (4)$$

where sp is the number of spare processors per tile.

Subelement-level redundancy is used within each processor as an additional means to improve cell (and wafer) yield. The Mosaic and PSDP processor designs each contain a RAM, which can include internal redundancy for yield enhancement. We conducted a preliminary study to estimate cell area and yield and to evaluate internal redundancy options. We found that the total area of the processor, excluding RAM, occupies less than 10 percent of the total chip area. By including redundant rows and columns in the RAM, we predict 90 percent RAM yields and processor cell yields of greater than 70 and 80 percent for Mosaic and PSDP, respectively.

Our yield estimates are based on a conservative defect density of $d = 3$ defects/cm^2 and a defect clustering parameter of $\alpha = 0.714$.

Programmable systolic data processor. A complete multicomputer system can be constructed using the PSDP processing element building block (see the sidebar). The homogeneity of this architecture maximizes the potential use of all functional cells on each wafer.

PSDP architecture for WSI

A joint effort is underway at Honeywell and CMR to study homogenous multicomputer designs for wafer-scale implementation. We are developing the PSDP architecture to achieve throughput advantages over its predecessors. Architectural innovations include distributed control and software allowing each processor to act as a self-controlled node. Details of

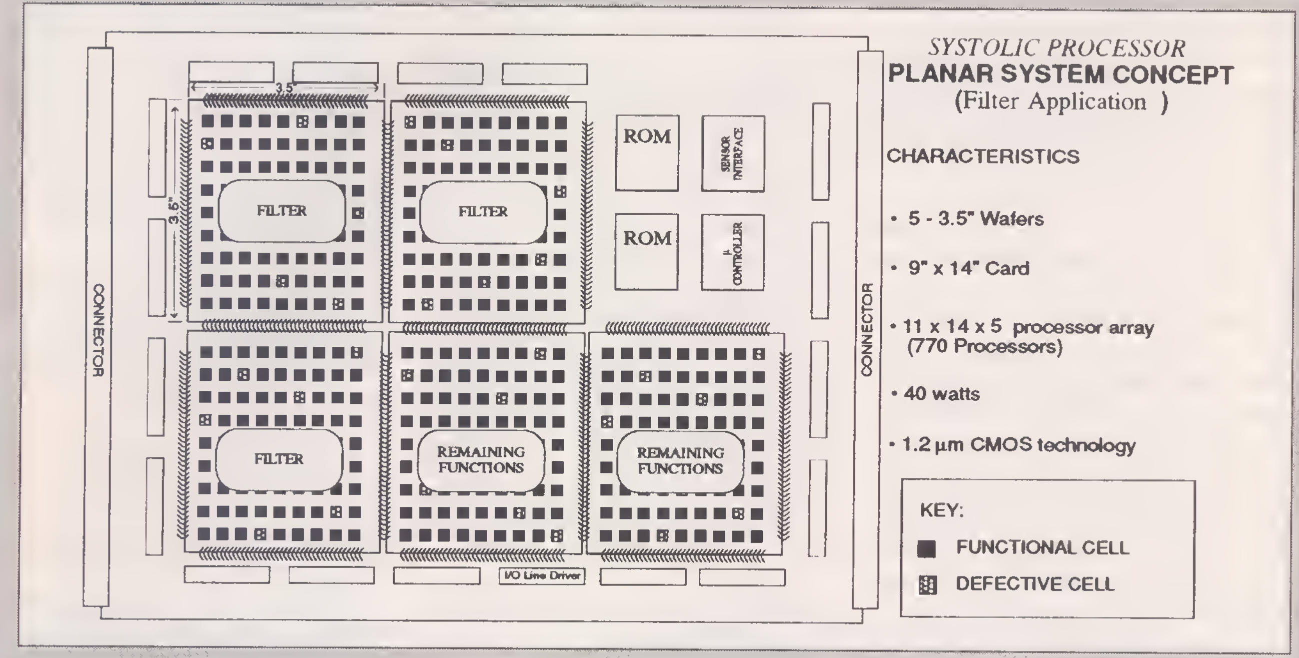

PSDP WSI prototype card.

Nearest neighbor communication (including virtual nearest neighbors) and low cell I/O count minimize the wafer area required for processor-to-processor interconnections. With a projected cell size of 5mm × 6mm (including interconnect wiring and power busing), 238 processors can be fabricated on a 5-inch wafer. As previously described, the predicted cell yield is more than 80 percent. However, studies have shown that the maximum two-dimensional mesh that can be configured (harvested) uses 70-75 percent of the good tiles. Thus, we can harvest only an 11 × 14 two-dimensional processor mesh per wafer.

Table 1 shows the computational availability predictions for the PSDP wafer. The first column indicates the tile organization: The first number identifies the number of working elements required per tile, and the second number indicates the total number of elements per tile. The next column identifies the tile yield associated with the particular tile design. The remaining columns indicate a measure of the system's computational availability (number of working processors) at various times.

Table 1. PSDP computational availability.

	Tile Yield	$t = 0$	$\lambda_0 t_0$	$2\lambda_0 t_0$	$3\lambda_0 t_0$	$4\lambda_0 t_0$	$6\lambda_0 t_0$
$c = .93$							
1/1	0.85	154	90.00	54.43	33.55	21.06	8.75
1/2	0.98	116	77.00	51.91	35.51	24.64	12.34
2/4	0.99	116	75.14	49.29	33.23	22.98	11.70
3/4	0.89	120	73.30	46.75	31.00	21.23	10.86
$c = .95$							
1/1	0.85	154	103.22	71.14	49.72	35.20	18.36
1/2	0.98	116	85.44	63.65	47.92	36.45	21.71
2/4	0.99	116	83.28	60.06	44.18	33.10	19.51
3/4	0.89	120	81.14	56.61	40.60	29.80	17.13
$c = .97$							
1/1	0.85	154	118.32	92.92	73.61	58.80	38.38
1/2	0.98	116	94.80	78.00	64.61	53.84	38.07
2/4	0.99	116	92.28	73.12	58.65	47.56	32.21
3/4	0.89	120	89.80	68.50	53.12	41.80	26.87

this WSI architecture and its physical implementation are presented in Landis et al.[1]

The PSDP uses a programmable systolic processing element that communicates with its neighboring cells over a 4-bit data path. This communication can be performed concurrently with internal 8-bit arithmetic logic unit (ALU) operations and 4 × 4-bit multiplications. The PSDP internal cell architecture is constructed from a few functional blocks: RAM, ALU, multiplier, latches, and multiplexers. The PSDP architecture takes advantage of an orthogonal memory structure that allows 8-bit data access as well as 32-bit instruction access. Distributed memory consists of a 2-Kbit dual-port static RAM located within each processor cell. The PSDP system can implement both single-instruction multiple-data (SIMD) and multiple-instruction multiple-data (MIMD) algorithms.

Cell size is estimated at 5mm × 6mm (including intercell wiring area) using 1.2-micron complementary metal-oxide semiconductor design rules. Optimization of the PSDP architecture for WSI includes power minimization, enhanced testability, clock distribution, as well as design optimization for defect tolerance and intercell connectivity.[1] The processor cell has a projected performance of 9 million 8-bit operations per second.

The figure shows a proposed planar prototype system of five wafers with each wafer containing 154 working processors.

Reference

1. D. Landis et al., "A Wafer-Scale Programmable Systolic Data Processor," *Proc. Ninth Biennial University/Government/Industry Microelectronics Symposium*, IEEE Press, Piscataway, N.J., 1991, pp. 252-256.

Failure rate is normalized with respect to time, so that the tabular data can be applied for different times (t_0) and processor failure rates (λ_0). The tabular data results from our choice of $\lambda_0 t_0 = 0.046$. This choice could correspond to a failure rate of $\lambda_0 = 5.24 \times 10^{-6}$ (failures/hour) at $t_0 = 1$ year, or to a failure rate of $\lambda_0 = 1.31 \times 10^{-6}$ (failures/hour) at $t_0 = 4$ years.

The tile design yielding the most computational capacity varies as the mission time increases. For $c = .93$, the (1/1) strategy is best up to $2\lambda t$; but by $3\lambda t$, the (1/2) design has higher capacity. By $4\lambda t$, the (1/1) approach has the least capacity of strategies evaluated. The table also indicates how the results change as coverage (c) is increased from 93 percent to 95 and 97 percent. When the coverage increases to 95 percent, the (1/1) is best up to $3\lambda t$; and for a 97 percent coverage, the (1/1) tile design is best over all mission times.

Mosaic multicomputer. Mesh-connected architectures such as the Mosaic multicomputer (see sidebar on p. 70) can be mapped efficiently to wafer scale. Assuming 5-inch wafers and approximately 9 mm × 9 mm processor cell area (including interconnect wiring and power busing), we project that a 10 × 10 array of processors can be fabricated on a single wafer. An estimated cell yield of 70 percent is predicted with subelement redundancy in the RAM and ROM. Hard restructuring is used to harvest the good processors. The high yield of the router, 99 percent, allows a choice of several different index mapping harvesting approaches to connect all 70 good processors per wafer.

Table 2 indicates that for this design, the (1/1) structure provides the highest computational availability for a coverage value of 97 and 95 percent. At 93 percent coverage, the (1/2) computational availability becomes slightly high-

Carnegie Mellon University
Software Engineering Institute

Career Opportunities

Shape a career.

We are looking for educators who have–

10 years of relevant teaching experience

Superior oral and written communication skills

Commitment to education as a mechanism for advancing software engineering practice

Technical experience in one or more areas of software engineering

Demonstrable leadership in educational innovation

Relevant advanced degree

We have openings for qualified applicants in the following areas–

Continuing education

Academic education

Send resumé to

Sally Miller
Software Engineering Institute
Carnegie Mellon University
Suite 414
Pittsburgh, PA 15213–3890

The Software Engineering Institute is an affirmative action, equal opportunity employer sponsored by the U.S. Department of Defense.

Table 2. Mosaic computational availability.

	Tile Yield	t=0	$\lambda_0 t_0$	$2\lambda_0 t_0$	$4\lambda_0 t_0$	$8\lambda_0 t_0$	$10\lambda_0 t_0$
c = .93							
1/1	0.70	70	53.74	41.64	25.66	10.72	7.23
1/2	0.91	45	37.39	31.25	22.19	11.90	8.95
2/4	0.92	45	36.61	29.46	19.67	9.69	7.08
3/4	0.65	48	36.46	28.17	17.54	7.72	5.36
c = .95							
1/1	0.70	70	57.20	47.06	32.46	16.55	12.18
1/2	0.91	45	38.90	33.79	25.78	15.69	12.49
2/4	0.92	45	38.06	31.75	22.59	12.29	9.33
3/4	0.65	48	37.92	30.32	19.99	9.52	6.78
c = .97							
1/1	0.70	70	60.88	53.17	41.03	25.47	20.45
1/2	0.91	45	40.42	36.52	29.94	20.66	17.37
2/4	0.92	45	39.57	34.21	25.92	15.54	12.24
3/4	0.65	48	39.43	32.63	22.76	11.70	8.55

er at long mission times. Finally, the effect of coverage is less than it was for the PSDP design, due to the smaller number of processors available on the Mosaic wafer.

Conclusions. These two WSI case studies illustrate how the optimal WSI design point varies with system architecture. Based on computational availability, the appropriate architecture varies with mission time/failure rate. For example, the (1/2) PSDP organization starts out with less computational capacity than the (1/1) approach but ends up with the more capacity as mission time or failure rate increases. The (1/1) approach may be preferable for a commercial computer whereas the (1/2) approach is better suited for a deep-space probe.

For very fine-grained architectures, there may be an advantage to placing multiple processing elements within a single tile. For less fine-grained architectures, it is more advisable to have a single processor per tile. This result also holds true from a yield perspective. For large cells with relatively low yield, the placement of multiple processing elements on a tile (even with spares) will produce a lower overall number of working elements per wafer.

Finally, system coverage affects the need for redundancy. The (1/2) and (2/4) tile structures have essentially the same number of spares and might be

Mosaic multicomputer architecture

Wafer-scale implementation needs fine-grain computing nodes, where each node is a relatively simple processor with tens of kilobytes of memory. A fine-grain multiprocessor becomes inefficient for large numbers of nodes because of the long network delay associated with accessing data from a global memory. However, this memory-access latency problem does not exist in fine-grain multicomputers with distributed memory. The Mosaic multicomputer architecture developed at Caltech[1] is such a fine grain multicomputer and makes an interesting case for wafer-scale implementation.

Several characteristics of the Mosaic multicomputer make it amenable to wafer-scale system implementation. The fundamental cell is a 16-bit RISC-like processor with a projected performance of 14 million instructions per second. A complete MIMD multicomputer system can be constructed from a single-cell microcomputer building block. The homogeneity of this architecture maximizes the potential use of all functional cells in a wafer-scale system. The low I/O count implicit in the byte-serial communication protocol minimizes the wafer area required for processor-to-processor interconnections.

The mesh-connected asynchronous message-passing communication network supports maximum performance from a wafer-scale system. It also creates a need for a new wafer-scale paradigm. The traditional concept of processing elements surrounded by programmable switches is not applicable to this design. The message-passing processor can be configured into a working system as long as all the routers are working. The runtime software can be changed to use only healthy processors, preventing any tasks from migrating to faulty processors. Therefore, reconfiguration requires no additional programmable switches.

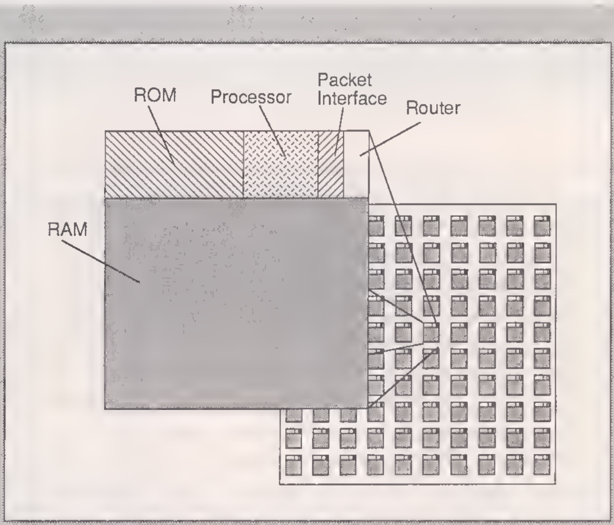

Mosaic multicomputer cell.

The figure shows an array of Mosaic processors, highlighting the key elements.

Reference

1. W.C. Athas and C.L. Seitz, "Multicomputers: Message Passing Concurrent Computers," *Computer*, Vol. 21, No. 8, Aug. 1988, pp. 9-24.

expected to perform in an identical fashion. However, two principal parameters come into play in the evaluation of computational availability: the coverage (c) and the failure rate (λ). The (1/2) organization's failure rate is only half that of the (2/4) since it has only half the number of operating processors. However, there are twice as many of these tiles. For lower values of coverage, the amount of computation available late in the system lifetime is nearly equal, but when the coverage is increased, the difference becomes much more prominent. In fact, the higher the coverage the less tile-level redundancy is needed. ■

References

1. W.C. Athas and C.L. Seitz, "Multicomputers: Message Passing Concurrent Computers," *Computer*, Vol. 21, No. 8, Aug. 1988, pp. 9-24.
2. M.D. Beaudry, "Performance-Related Reliability Measures for Computing Systems," *IEEE Trans. Computers*, Vol C-27, No. 6, June 1978, pp. 540-547.

David L. Landis is an associate professor in electrical engineering at the University of South Florida and a research scientist in the Center for Microelectronics Research there. His research interests are fault-tolerant processing, design for test, and VLSI/WSI design.

Nitin Nigam is working toward his PhD in electrical engineering from the University of South Florida. His research interests include fault-tolerant VLSI systems, built-in self-test, and VHDL modeling.

Joseph W. Yoder is a research associate in the Center for Microelectronics Research at the University of South Florida, where he is working toward his PhD in electrical engineering. His research interests are in VLSI/WSI design, including fault-tolerant design, design for test, and rapid prototyping.

An Architecture for WSI Rapid Prototyping

Vijay K. Jain, Hiroomi Hikawa, and D.C. Keezer
University of South Florida, Center for Microelectronics Research, Tampa, FL 33620-5350

Monolithic wafer-scale circuits can provide significant gains in throughput and reliability over more conventional circuits, while simultaneously reducing weight, volume, and power requirements. Recent progress in wafer-scale integration (WSI) has come through synergistic advances in several areas, including semiconductor processing and the technology for interconnections. One of the most important, however, is in recognizing defects on the wafer and circumventing them through redundancy and reconfiguration. In fact, without redundant components and a reconfiguration strategy for replacing defective components, the yield of a wafer-scale device tends to zero.

Significant contributions in this area have been made by Sami, Stefanelli, Saucier, Raffel, Wyatt, Fuchs, Kung, and Singh, among others.[1] To facilitate this defect-circumvention process, as few types of cells as possible must be used and the cell I/O simplified. Further, statistical considerations warrant that the cells themselves be considerably smaller than state-of-the-art VLSI chips. At the wafer level, however, performance much above the state of the art for a single chip can be achieved.

Signal processing algorithms possess a high degree of regularity. Thus, they are suitable for WSI because they facilitate the mapping of the algorithm to an array of just a few types of cells on the wafer. During the past two years we have conducted research on a radix-8 WSI fast Fourier transform (FFT) processor.[2] The goal has been to design a word-parallel, multiprocessor architecture with high throughput. At the board level, the objective was to ensure that this radix-8 wafer design can serve as a workhorse for high-speed, 512- or 4,096-point frame FFTs.

We also developed a dedicated triangular lower-upper (L-U) matrix decomposition architecture. (Note that this algorithm is being implemented in a dedicated wafer design at Lincoln Laboratory.[1]) As in the case of the FFT algorithm, the L-U decomposition algorithm is computation intensive and highly regular. But unlike the FFT, it requires a nonlinear operation, namely, the computation of the reciprocal of certain matrix elements.

This article presents a generalized architecture for signal processing — a WSI architecture for rapid prototyping, or WARP. Rapid prototyping aims at using a single-wafer architecture to map one of several members of a class of algorithms. The wafer can be personalized for the algorithm by either soft or hard (laser) restructuring. The WARP wafer consists of an array of two types of cells specifically defined for this architecture. They are the universal multiply-subtract-add (UMSA) cell and the universal nonlinear (UNL) cell.

In the past, signal processing WSI and ultra large scale integrated (ULSI) architectures primarily addressed multiply-accumulate operations, applicable to finite-impulse response (FIR) filtering and FFT computation. The WARP architecture attempts to broaden the class of algorithms by using a single rapid-prototyping architecture. This new architecture can implement not only the FFT computation and L-U decomposition algorithms but also algorithms such as FIR filtering, Cholesky decomposition, and linear system solution. Recently, we have also mapped such advanced signal processing algorithms as the singular value decomposition (SVD) algorithm to WARP.

WARP. Here we describe the UMSA

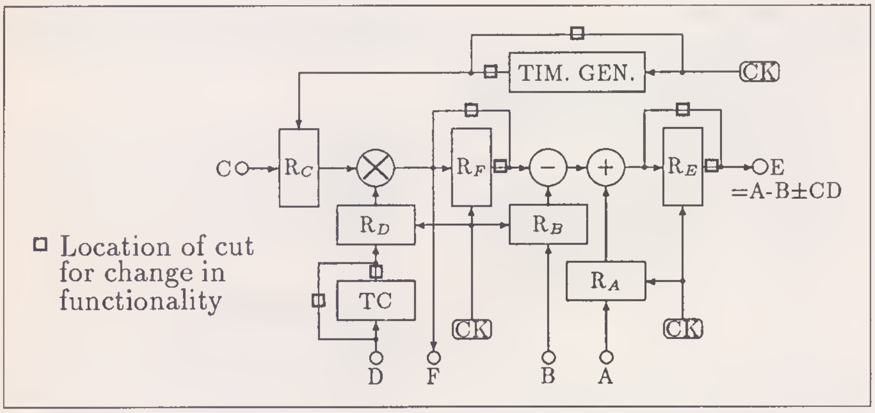

Figure 1. UMSA architecture.

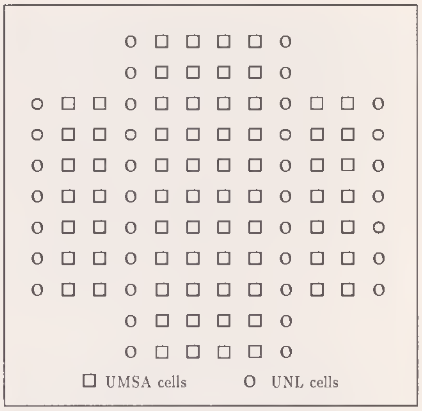

Figure 3. WARP wafer floor plan.

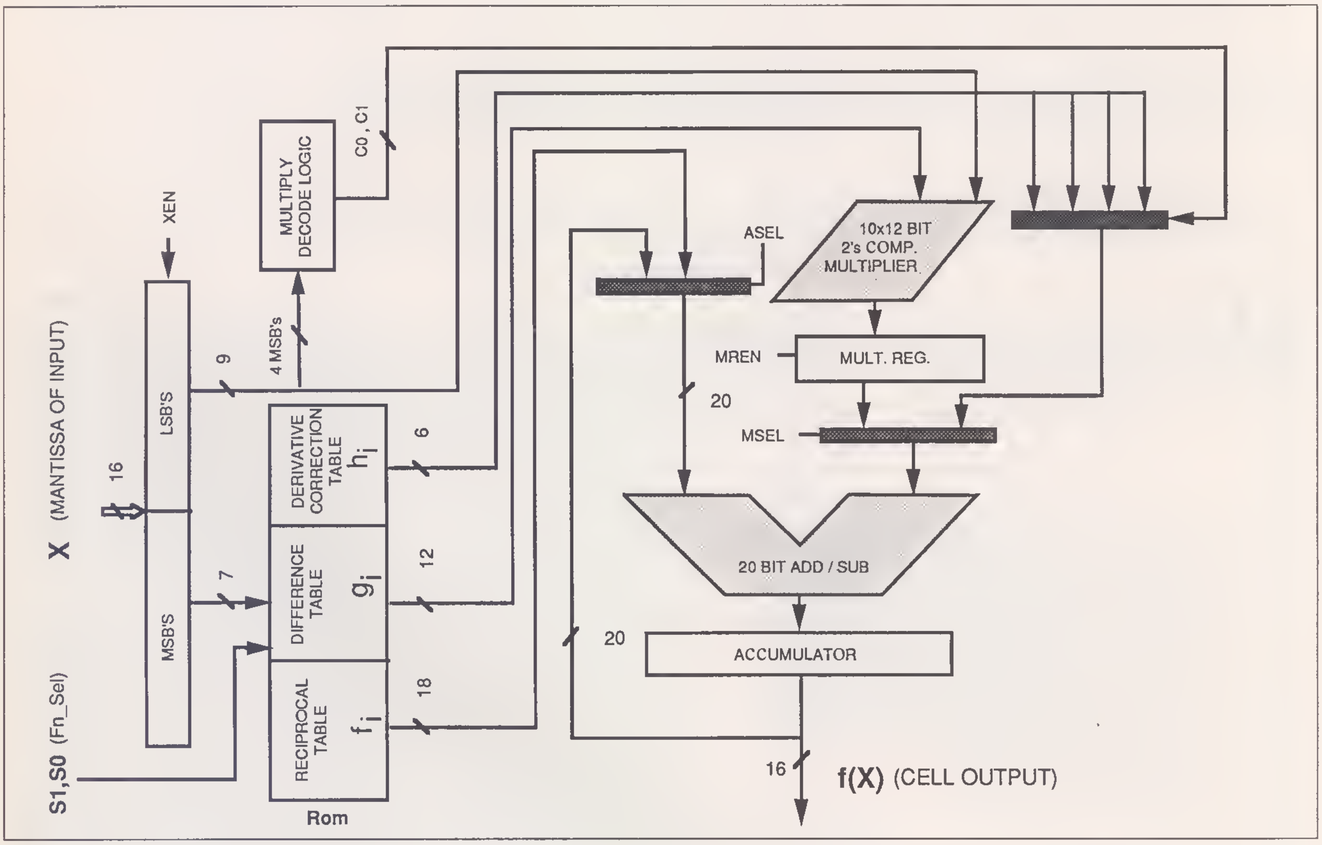

Figure 2. UNL architecture.

and UNL cells and the mapping of the FFT and L-U decomposition algorithms to WARP.

Universal cells. The key to WARP is the definition of the high-speed universal cells. The adjective "universal" is, of course, relative to the domain of algorithms planned. The domain for our effort was conceived as discussed above. The UMSA and UNL cells devised for this purpose are presently designed for 16-bit parallel arithmetic. The UMSA consists of a multiplier and two adder/subtracter units, as shown in Figure 1. Note that the cell's functionality can be changed by performing cuts at the locations indicated in the figure.

Defining a UNL element for the targeted applications required a significant amount of research.[3] Figure 2 shows the UNL architecture. The cell generates one of several functions (reciprocal, square-root, sine/cosine, or arctangent) in two clock cycles after the pipeline is filled. Signal delay estimates for 2-micrometer complementary metallic oxide semiconductor and 1.2-micrometer CMOS are 120 and 80 nanoseconds, respectively.

Table 1. Results of exhaustive runs for 16-bit UNL.

Function	Normalized Range	Input Format	Output Format	Number of Tests	Least-Significant-Bit Errors	Average Error
Reciprocal	$0.5 \le x < 1$	(0,16)	(2,14)	32,768	3,725	0.0000068
Square Root	$0.25 \le x < 1$	(0,16)	(0,16)	49,152	4,235	0.0000012
Sine Function	$0 \le x < \pi/2$	(1,15)	(1,15)	65,536	12,102	0.0000021
Arctangent	$0 \le x < 1$	(0,16)	(1,15)	65.536	3,544	0.0000021

Table 1 summarizes the UNL cell's arithmetic performance. The cell is used for sine/cosine generation in FFT computations. In L-U decompositions, it computes the reciprocal function. In fact, for the Cholesky decomposition algorithm, the UNL is used both for the reciprocal and square-root operations. Jain et al.[3] present details pertaining to the implementation of this algorithm.

Rapid-prototyping architecture. Figure 3 shows the WARP floor plan. It is based on an array of 72 UMSA cells and 36 UNL cells. The wiring channels are used for both inter-UMSA as well as UMSA-UNL connections.

Figure 4 shows the mapping of the FFT butterfly, which can be implemented with four UMSA cells (in the presence of defects). Similarly, the L-U decomposition algorithm can be mapped to the WARP architecture. For simplicity, Figure 5 shows a mapping for a 4 × 4 matrix. The connections are somewhat similar for an 8 × 8 matrix.

Reconfiguration. Here we discuss reconfiguration of the algorithms in the presence of defects.

FFT algorithm. We have developed a computer program that reconfigures the wafer in the presence of defective cells. It uses the concept of pooled redundancy to form blocks, each containing enough resources to realize a butterfly. A random-number generator produces a defect map in the simulation mode. Figure 6 summarizes the algorithm. It is heuristic and based upon borrowing from the next higher or next lower row of cells. If the reconfiguration is successful, the result is a connection map. Figure 7 presents an example.

L-U decomposition algorithm. We have also developed a computer program to perform reconfiguration for the L-U decomposition algorithm. In contrast to the FFT case, the goal here is to complete one row of matrix manipulation at a time. Starting with the top row of the 8 × 8 matrix, the reconfiguration progresses through all rows of the matrix operations. Figure 8 summarizes the details.

Harvesting probability model and yield. A homogeneous wafer is a wafer with only one type of cell. If the wafer is provisioned with N cells and the target design needs M cells, then the wafer yield is given by

$$Y_{\text{wafer}} = \sum_{k=M}^{N} p_k h_k \qquad (1)$$

where P_k is the binomial probability of having exactly k healthy cells, and h_k is the probability of configuring a functional wafer given exactly k healthy cells.

It is assumed that the yield of the interconnect and the switches is high enough for the cell yield

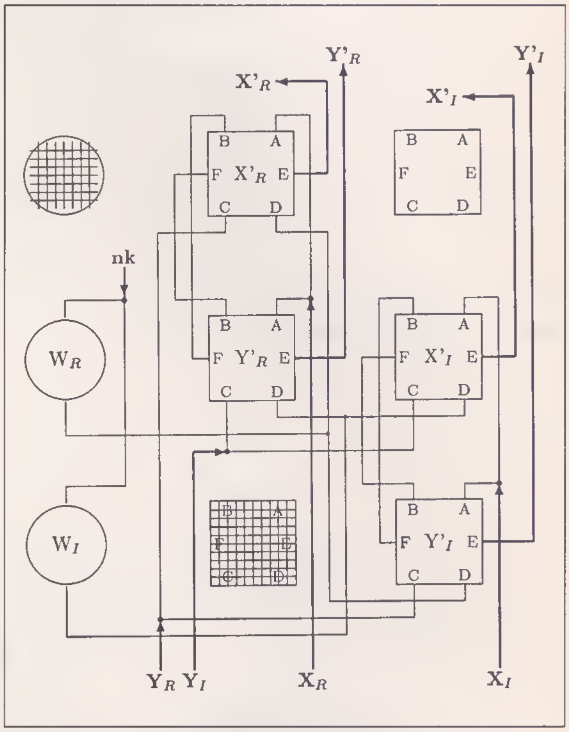

Figure 4. Butterfly interconnections.

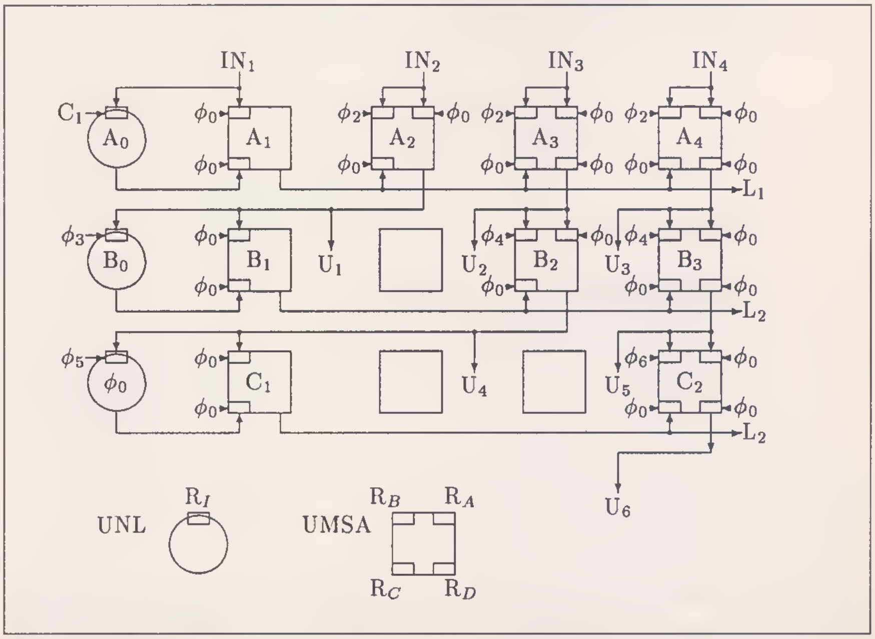

Figure 5. Interconnections for a 4 × 4 L-U decomposition.

```
FFT Reconfiguration:    Preassign all Butterflies;
                        NBUTTERFLY = 0;
Next Butterfly:         if end of Butteryfly; goto Exit;
                        NDEFECT = number of Defective Cells in the Butterfly;
                        if NDEFECT = = 0: goto Success;
Next Cell:              search for HEALTHYCELL in spare cell area;
                        if no HEALTHYCELL found: goto Next Butterfly;
                        if the HEALTHYCELL mingles with other Butterflies:
                        goto Next Cell;
                        replace the DEFECTIVECELL with the HEALTHYCELL;
                        NDEFECT = NDEFECT – 1;
                        if NDEFECT = = 0: goto Success;
                        goto Next Cell;
Success:                NBUTTERFLY = NBUTTERFLY + 1;
                        goto Next Butterfly;
Exit:                   if NBUTTERFLY = 12: Reconfiguration success;
                        else: Reconfiguration failed;
```

Figure 6. FFT reconfiguration algorithm.

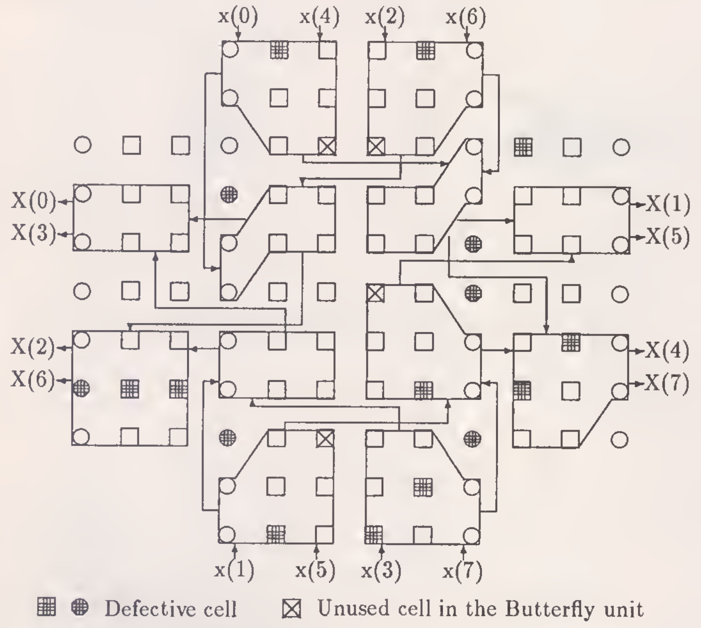

Figure 7. FFT reconfiguration on a WARP wafer.

to predominate. The critical information needed, then, is the harvesting probability function h_k. We model this probability function, using a flexible model that appears to fit the experimental harvesting data obtained by a number of investigators. This model is described by

$$h_k = \frac{c+1}{c+\left(\frac{(L-k)}{(L-N)}\right)^2} \qquad e_k = r_k e_k \tag{2}$$

where e_k is the exponential function described by Jain, Hikawa, and Swartzlander.[1] The factor r_k has been introduced in Equation 2 to match the experimental data at the inflection point. The procedure for determining parameters a, b, c, and L is discussed in Jain, Keezer, and Hikawa.[4]

A computer program was written to perform the harvesting simulations using the pooled redundancy concept. A block containing at least four good UMSAs and two UNLs was used to configure a butterfly, as discussed in the previous section. Harvesting data was obtained from 10,000 runs, each of which used a random generator to create a "wafer defect map." The square symbols in Figure 9 represent this data.

The 8-point FFT function requires 48 functional UMSAs, while the wafer has $N = 72$ cells.[1] We estimated the inflection point at $L = 65$, for which the simulation-generated harvesting data gives $h_L = 0.655$. The other point chosen for estimation is h_{58} (that is, $k = 58$). The value of c was found to be $c = 0.85$. The solid line in Figure 9 shows the resulting harvesting model and has the parameters $a = 0.09$, $b = 1.0$, $c = 0.85$, and $L = 65$.

Using the model values in Equation 1, $Y_{\text{wafer}} = 0.613$. We used Equation 2 with $p_{\text{cell}} = 0.9$ to compute p_k (shown by the dashed curve in Figure 9, magnified five times). Actually, this should be denoted as $Y_{\text{wafer,MSA}}$, since it is based on multiply-subtract-add cells alone. A similar analysis gives $Y_{\text{wafer,ROM}} = 0.85$, so that $Y_{\text{wafer}} = (0.613)(0.852) = 0.52$.

Wafer-scale testing and test facilities. Testing WSI circuits presents challenges beyond those encountered in VLSI. Careful test planning and strict adherence to design-for-test practices are of paramount importance to the success of the WSI system. Boundary scan, internal scan, and built-in self-test represent some common approaches that are used to alleviate an otherwise intractable testing problem. Clearly, the first step toward harvesting is to obtain a defect map. Normally, this is accomplished with an automated electrical test system connected to the cell through a "probe card" interface.

A probe card provides a means of contacting the small (typically 75- to 100-micrometer) pads of the cell. If the cell has relatively few I/O pads, then a probe card can be designed to connect to all these pads. In some cases, the wafer area required for probing can be minimized by adopting a standard (small pin-count) common test port working together with boundary scan and possibly internal scan registers. However, regardless of the number of I/O lines supported by the cell, a die sort using a probe card is generally performed to identify the functional cells.

The University of South Florida test system uses a Hewlett Packard model 82000-D200 tester to support devices with up to 200 bidirectional channels, operating at rates up to 200 MHz. The system achieves overall accuracy of ±250 picoseconds using a 50-picosecond edge-placement resolution available on each channel. A Thermonics model T-2420 temperature forcing system controls the temperature of individual devices during testing. Further, an Electroglas 2001x wafer prober supports the wafer probe card and manipulates the wafer itself. It includes the optics for observing the fine alignment required between the wafer and the probe card. The same equipment is used to test the completed WSI system after reconfiguration. ■

```
8 × 8 L-U
Reconfiguration:   NBLOCK = 0; Row = 3;
Next BLOCK:        if end of BLOCK: goto Exit;
                   MCELL = number of cells required in the BLOCK;
                   NCELL = 0;
Search:            Search HEALTHYCELL in Row(Current row);
                   if no more HEALTHYCELL in ROW: goto Upper Search;
                   NCELL = NCELL + 1;
                   if NCELL = = MCELL: goto Success;
                   goto Search;
Upper Search:      Search HEALTHYCELL in ROW – 1 (Upper row);
                   if no more HEALTHYCELL in ROW – 1: goto Lower Search;
                   NCELL = NCELL + 1;
                   if NCELL = = MCELL: goto Success;
                   goto Upper Search;
Lower Search:      Search HEALTHYCELLs in ROW + 1 (Upper row);
                   if no more HEALTHYCELL in ROW + 1: goto Fail
                   NCELL = NCELL + 1
                   if NCELL = = MCELL: goto Success;
                   goto Lower Search;
Success:           NBLOCK = NBLOCK + 1;
Fail:              ROW = ROW + 1;
                   goto Next BLOCK;
Exit:              if NBLOCK = = 7: Reconfiguration success;
                   else: Reconfiguration failed;
```

Figure 8. Reconfiguration for 8 × 8 L-U decomposition.

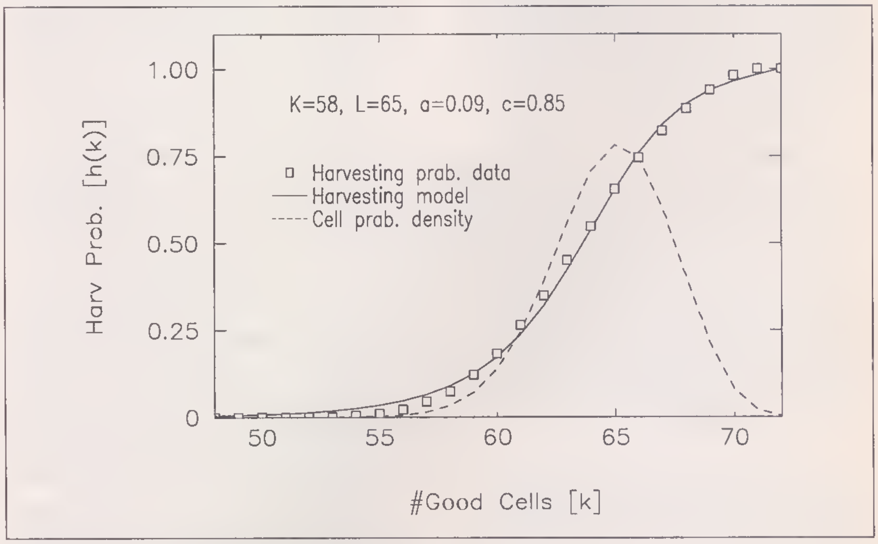

Figure 9. New barvesting probability model.

Acknowledgments

This research was supported in part by National Science Foundation Grant No. MIP-9103286, by the Florida High Technology and Industry Council, and by DARPA Grant No. MDA 972-88-J-1006.

The authors are deeply indebted to Earl E. Swartzlander Jr. for his helpful comments on this research.

References

1. *Proc. Int'l Conf. WSI*, M. Little and V.K. Jain, eds., IEEE CS Press, Los Alamitos, Calif., Order No. 2126, 1991.
2. V.K. Jain et al, "Fully Parallel Wafer-Scale Architecture for an FFT Processor", *Int'l J. Computer-Aided VLSI Design*, Vol. 3, No. 1, 1991, pp. 113-135.
3. V.K. Jain, G.E. Perez, and J.M. Wills, "Novel Reciprocal and Square-Root VLSI Cell: Architecture and Application to Signal Processing," *Proc. Int'l Conf. Acoustics, Speech and Signal Processing*, IEEE Press, Piscataway, N.J., 1990.
4. V.K. Jain, D.C. Keezer, and H. Hikawa, "Defect Tolerance and Yield for a WSI Rapid Prototyping Architecture," *Proc. Int'l Workshop Defect and Fault Tolerance in VLSI systems*, IEEE CS Press, Los Alamitos, Calif., Order No. 2457, 1991, pp. 24-27

Vijay K. Jain is a professor of electrical engineering at the University of South Florida. His research interests include VLSI/WSI design and applications, and high-speed signal processing systems.

Hiroomi Hikawa is a postdoctoral researcher at the University of South Florida. His research interests include parallel processing algorithms, architecture for communications and signal processing, and VLSI/WSI implementations.

D.C. Keezer is an associate professor of electrical engineering at the University of South Florida and the coordinator of the Microelectronics Test Laboratory at USF's Center for Microelectronics Research. His research interests include VLSI/WSI design and test, design for test, and high-speed and mixed-signal testing.

Computer Society NEWS

Editor: Guylaine M. Pollock, Sandia National Laboratories, Division 1424, PO Box 5800, Albuquerque, NM 87185; phone, (505) 845-7463; Internet, gmpollo@cs.sandia.gov

NEWS FROM THE COMMITTEE ON PUBLIC POLICY

Development to begin on standards of ethical practices for computing professionals

Robert J. Melford, Chair, COPP Computing Ethics Subcommittee

Casually scanning issues of our profession's technical journals — *Computer*, *Communications of the ACM*, *IEEE Spectrum*, and others — published over the past few years, we find many references to escalating concerns about ethics and our professional activities. Some have been dramatic, such as John P. Barlow's call for "the development of a new social contract for the digital domain — sort of a cyberspace equivalent to the Code of the West — and an attendant definition of the rights and responsibilities of the inhabitants."[1] Others, like Michael McFarland's, have been academic and comprehensive:

> Establishing a set of [ethical] standards that is accepted and shared by the whole profession is important not only because it can help those facing difficult decisions to make reasonable and fair judgements that serve the best interests of everyone involved, but also because of the sense of support and solidarity such standards can give those making the decisions, making it easier for them to carry through on those decisions with courage and confidence.[2]

Governmental, industrial, and academic institutions have begun focusing considerable attention on ethics and computing:

- The National Science Foundation has funded the development of books and the presentation of international conferences on computer ethics.
- The Computer Science Accreditation Board now recommends that the study of ethics be included in the computer science curriculum. The Association for Computing Machinery's curriculum task force has suggested similar guidance.
- Consulting firms specializing in ethical issues have begun to appear.
- The New Jersey state legislature has considered a bill to require the licensing of all software designers.
- The president of the ACM has suggested studying the possibility of licensing all computing professionals.
- The IEEE Computer Society's Committee on Public Policy (COPP) is reviewing the advisability of drafting an "Entity Position Statement on Software Certification by Licensed Software Engineers" for possible consideration by the society's Board of Governors.

These activities largely address problems of computer misuse and other computer security matters. Books, articles, and conference presentations on computing ethics often amount to little more than an inventory of the ways people have committed crimes with computers. They may discuss how individuals have used public- and private-sector computing resources for their own financial gain, or explore some of the potential second-order consequences to civil rights and freedoms. Examples range from how the potential consolidation of distributed databases may threaten an individual's sense of anonymity and privacy, to the problems stemming from erroneous information stored in credit or other databases.

However, ethics for and within the computing profession is a different topic altogether from the use of computing resources in the violation of more general ethical principles. For an analogy, let's recall the misconception still shared by many congressional staffers and other policymakers who believe that by setting up the four NSF supercomputer centers they have given vast financial and technical resources to computer science research. They forget that these centers are used primarily by scientists in other disciplines. Like many people, they associated the funding of computers with the funding of computer science and engineering.[3] Likewise, past ethical exercises have addressed issues tangential to the major ethical issues faced by practicing software engineers.[2,4,5]

Ethical issues for computing professionals are a consequence of our technical specialty. What does the software engineer do when management insists on shipping safety-critical software because the contractual delivery date has arrived, yet the engineer still has professional doubts about the software's ability to function correctly when stressed in its intended real-world applications environment? What responsibilities does the computing profession have to support the engineer with resources to adjudicate this conflict? And what about the ethical issues of the next century: Will we have built the necessary foundations for addressing them?

These and other questions yet to be asked can only be answered when the computing profession, after carefully reflecting on experience and considering the relative importance of conflicting values, formulates norms and guidelines that will help individuals identify and subsequently avoid, minimize, or resolve ethical conflicts. Such norms emerge through a long iterative process of argument, testing, and revision. As their soundness and workability are proved through application, they mature to become "widely accepted as binding on those subject to them."[2] Historically, this has been true of the professional norms defining the obligations of doctors to patients, or of attorneys to their clients and the judicial system.

The IEEE standards process provides an ideal vehicle for identifying, discussing, and defining appropriate norms and guidelines in a structured and coordinated fashion. While the lengthy and deliberative process of developing the profession's consensus on a given standard has often been criticized for failing to keep pace with changing technologies, in the area of ethics this can be advantageous. Standards of professional ethical conduct must have a quality of timelessness, for once they are found valid and applicable, their value does not abate over time.

We must, of course, recognize that as engineers and technologists we are generally uncomfortable with ethical issues,[2,5,6] choosing instead to apply ourselves only to areas within our sphere of perceived competency. We gravitate toward areas amenable to rigorous analysis — namely, technical problems — on the basis of our preexisting and well-defined knowledge and experience. But we should not let this deter us from accepting the challenge. As Barlow pointed out,

> The trouble is, the rest of society is so utterly perplexed by digital technology that most ordinary citizens are even less qualified than [computer scientists and engineers] to engage in cybernetic social philosophy. . . . You are almost certainly more knowledgeable about the legal and cultural ambiguities surrounding digital [computing technology] . . . than your computer-phobic fellow citizens.[1]

And McFarland warned, "If they [computer professionals] do not act on their own, they will find themselves subject to legislation made by people who do not really understand the issues involved."[2]

So there we have our incentive to move forward. But what about the skills, knowledge, and organization required? We *do* have relevant experience. In addition to promoting and facilitating the development of technical standards, the IEEE and the Computer Society have had considerable experience with managerial (as it relates to software production) and academic standards. The IEEE has promoted the development and application of a professional code of ethics for engineers.

Currently approved IEEE standards publications fall into three classes:[7]

- *Standards:* These are documents with mandatory requirements (generally characterized by use of the verb "shall").
- *Recommended practices:* These documents present the positions and procedures preferred by the IEEE (and normally use the word "should").
- *Guides:* Alternative approaches to good practices are suggested in these documents, but no clear-cut recommendations are made.

We can reasonably assume that most standards of ethical practices would be articulated as "recommended practices" and "guides."

The IEEE and the Computer Society membership includes many of the most highly accomplished, experienced, and respected professionals and academicians in computing science and engineering. Their technical expertise and relative objectivity have served the standards process admirably.

Like the physical sciences, ethics too has its own discipline and methodology. In McFarland's words, "Ethical knowledge emerges [in much the same way as knowledge in the physical sciences] from experience and reason, from action, and from reflection on that action."[2]

Most IEEE standards — for example, 802 or 1003 (Posix) — have been built on prior work. In the case of 802, the Digital/Intel/Xerox Ethernet communications protocol served as the basis for much of the early work. The current lack of a codified set of comprehensive and formally published recommendations on professional practices, procedures, and positions will necessarily constrain all efforts to formulate formal IEEE standards of ethics.

Because we lack this enabling resource, the Computing Ethics Subcommittee of the Committee on Public Policy proposed that the society conduct one or more workshops to prepare an informational document to support the establishment of a formal set of ethical standards. The proposal suggested that an interdisciplinary approach, drawing on a broad spectrum of professional and social expertise, would best lead to a focused discussion within the computing profession about the nature and scope of appropriate standards of ethical practices. The Computing Ethics Subcommittee recommended that representatives from the practicing and academic computing communities, as well as members of the ethical, legal/judicial, educational, ecumenical, business, and end-user communities, be encouraged to participate.

The first IEEE Computer Society Workshop on Ethical Standards for the Profession will be held May 8-9 in Washington, D.C. The workshop's preliminary session, Friday afternoon, will begin with a briefing on the IEEE standards history, process, and basic requirements. Such issues as the perceived nontechnical nature of ethics as opposed to the perceived technical nature of IEEE standards will be explored. Participants will also have a chance to identify and briefly discuss other issues affecting the development of ethical standards for the profession. Friday's program will conclude with an attempt to outline the areas of most immediate concern for the application of ethical standards. These are expected to include the software engineering cycle, safety-critical software, ethical practices in unethical environments,[2] forums for discussing and investigating ethical problems,[2] responsibilities of professional societies toward members, and conflicts of interest in formal standards development processes between the needs of the standard and the interests of participants' companies.

Saturday morning's session will concentrate on identifying and summarizing the various ethical issues raised in each of these subject areas. Depending on the number of participants, the workshop will either break into parallel task forces, one per subject area, or proceed sequentially through the subject areas as a single group. A task group could be formed to identify new subject areas, or this could be accomplished at the conclusion of the sequential discussion. The afternoon will be used to prioritize standardization activities on the basis of conclusions reached by the group or the task forces. The workshop will conclude with recommendations for subsequent action.

Conference and workshop information

For contact information on the Computer Ethics Institute's First National Conference, see the listing in this issue's Calendar section (May 7-8).

For additional information about the first IEEE Computer Society Workshop on Ethical Standards for the Profession, May 8-9, contact Ethical Standards Workshop, IEEE Computer Society, 1730 Massachusetts Ave, NW, Washington, DC 20036-1903, phone (202) 371-1013, fax (202) 728-9614; or Robert J. Melford, R.J. Melford & Associates, 24571 Tarazona, Mission Viejo, CA 92692-2350, phone (714) 830-1443, e-mail r.melford@ieee.org

The workshop will begin after the lunchtime conclusion of the Computer Ethics Institute's First National Conference, May 7-8, at the Brookings Institution, Washington, D.C. The theme of this conference, conducted in cooperation with COPP, is "In Pursuit of a 'Ten Commandments' for Computer Ethics." The conference will address three main topics. The first will be "Exploring the Range of Computer Ethics: Public Awareness, Public Policy, and Public Principles." The Thursday afternoon session will concern "Merging Theory and Practice: Identifying Organizational Mechanisms for Applying Ethical Principles and Making Moral Claims." The third session, Friday morning, will consider a "Blueprint for Action."

COPP's Computing Ethics Subcommittee expects to submit an IEEE Standards Project Authorization Request application in June for the development of ethical standards for computing professionals. Submission of this application will depend on a favorable outcome of the society's ethics workshop and the interest of the society's Standards Coordinating Committee in sponsoring the appropriate working group.

References

1. J.P. Barlow, "Coming into the Country," *Comm. ACM*, Vol. 34, No. 3, Mar. 1991, pp. 19-21.
2. M.C. McFarland, "Urgency of Ethical Standards Intensifies in Computer Community," *Computer*, Vol. 23, No. 3, Mar. 1990, pp. 77-81.
3. *1988 Snowbird Report*, D. Giers, chair, Computing Research Board, T. Walker, chair, Snowbird 88, "A Discipline Matures," *Computer*, Vol. 22, No. 2, Feb. 1989, pp. 72-75.
4. D. Gotterbarn, "Professional Software Engineering Ethics Education," *IEEE Computer Soc. Tech. Committee Privacy and Security Newsletter*, Oct. 1991, p. 5.
5. M.C. McFarland, "Ethics and the Safety of Computer Systems," *Computer*, Vol. 24, No. 2, Feb. 1991, pp. 72-75.
6. R.J. Melford, "Computer Ethics: The Technologists' Viewpoint," *Proc. Conf. Will Computers Be Ethical?* Brookings Institution, Washington, D.C., to be published.
7. *IEEE Standards Manual*, Inst. Electrical and Electronics Engineers, New York, Apr. 1988, p. 2.

Contest seeking computer aids for the disabled has 11 winners

A flexible keyboard, a musical ensemble program, and an audio-enabled program that allows blind users to interface with graphic computer programs won the top three awards in the Johns Hopkins national search for computer aids for the disabled. The Smithsonian Institution in Washington, D.C., hosted the final judging, along with a public exhibit.

The university's national search attracted 800 entries exhibited in 10 locations across the country, from which 30 regional winners were chosen to compete in the finals.

Top three prizes. Arjan S. Khalsa of Richmond, California, won first prize for his Unicorn Smart Keyboard, a membrane keyboard that recognizes large, colorful plastic overlays through which users access the computer by a simple touch of the hand.

Jonathon Adams from Cambridge, Massachusetts, placed second with his Switch Ensemble music program, which produces increasingly complex musical tunes when a user simply presses pads or skims screens.

Third place went to Frank McKiel from Trophy Club, Texas, for his audio-enabled graphic user interface, which allows blind persons to access Windows 3.0 programs and other graphical user interfaces.

Additional winners. The other inventions that won national awards are

- Tonetalker, a portable hand-held device that lets hearing-impaired persons see messages sent over a touch-tone telephone;
- the Peabody multimedia adult literacy program that tutors students through a digitized human voice on a videodisc;
- Read My Lips, a system to aid speech-reading for the hearing impaired using an interactive-media approach;
- Tiny Talk, a memory-resident program for the blind that translates

香 港 大 學

THE UNIVERSITY OF HONG KONG

Chair of Computer Engineering

(Ref. RF/TA/Engg/1)

Applications are invited for the Chair of Computer Engineering, a joint appointment in the Departments of Electrical and Electronic Engineering and Computer Science, tenable from 1 January 1993.

Applicants should have appropriate academic and professional qualifications and experience in a relevant field. They are expected to have a distinguished record in research and scholarly publications, demonstrated leadership and a strong commitment to teaching and curriculum development in Computer Engineering at both the postgraduate and undergraduate levels.

The University would prefer to make a permanent appointment, but consideration may also be given to applications for appointment on fixed or secondment terms of preferably not less than three academic years.

The University reserves the right not to fill the Chair or to fill the Chair by invitation or to make an appointment at a lower level.

Annual salary (superannuable) will be within the professorial range, of which the minimum is HK$685,620 and the average is HK$848,040 (approx. £49,680 and £61,450 respectively; sterling equivalent as at 18 February 1992). At current rates, salaries tax will not exceed 15% of gross income. Housing at a charge of 7.5% of salary, children's education allowances, leave and medical benefits are provided.

Further particulars and application forms may be obtained from Appointments (40259), Association of Commonwealth Universities, 36 Gordon Square, London WCIH OPF, UK; or from the Appointments Unit, Registry, The University of Hong Kong, Hong Kong (Fax (852) 5592058, E-mail APPTUNIT@HKUVM1.HKU.HK). Closes: *1 May 1992*.

modern software displays such as pop-up windows into speech;

- a program that uses graphic symbols in voice recognition to allow people with mental retardation or severe handicaps to communicate with others;
- a device to aid the disabled in controlling personal computers using a variety of switches, joysticks, and expanded keyboards;
- the SAM display device that allows nonverbal people to communicate using Morse code through a single switch; and
- LessonMaker, a talking interactive tool that lets teachers customize lessons for students' changing needs.

Funding. Paul Hazan, project director, organized the Johns Hopkins search with funds from the National Science Foundation, MCI Communications, and Microsoft. Inquiries should be addressed to Hazan at Johns Hopkins University, Applied Physics Lab, 1 W 128, Laurel, MD 20723.

Melbourne University takes top prize in programming contest

Melbourne University, Australia, beat 29 other finalists in the 16th annual Association for Computing Machinery Scholastic Programming Contest. Michigan State University and Stanford University finished second and third, respectively.

The five-hour battle of logic, strategy, and mental endurance brought together university teams from the US, Europe, Canada, and the Pacific Rim. ACM challenged each team to create programs to solve as many real-life puzzles as possible in the shortest time and with the fewest attempts.

The challenges, ranging from the serious to the whimsical, included

- creating a program to play the game Othello,
- charting a boat's course at sea through a series of radio signals,
- creating a program to eradicate moths in the Northeast, and
- determining which buildings in a city are visible from a southern elevation.

Each team was provided with a printer and a microcomputer using Pascal and C. Following a successful pilot run in last year's contest, the 1992 match used AT&T's Unix System V, which allowed contestants to communicate in their choice of Pascal or C. Once a team was ready to submit a program for evaluation, an AT&T StarLAN 10 Unix-based network moved the file to the appropriate judge for review.

The contest, held in Kansas City, Missouri, was sponsored by AT&T EasyLink Services because, in the words of the firm's president, Gordon Bridge, ". . . computer literacy and global messaging are becoming the basic tools of international commerce."

Applications available for 1992 Loebner Prize Competition

The quest for a thinking computer will shift into high gear November 17 when contestants take their systems to Boston for the 1992 Loebner Prize Competition.

The contest is a restricted version of the classic Turing Test of machine intelligence, named after the late British mathematician Alan Turing. If a computer ever succeeds at making judges believe it's a human being in an open-ended test, at least $100,000 will be awarded to the designers of the system and the contest will be abolished. This year's winner in the limited test will receive $2,000.

Last year's winner of the first annual Loebner Prize Competition was Joseph Weintraub, president of Thinking Software, Woodside, New York, for his computer program PC Therapist, which operated in the domain of "whimsical conversation." PC Therapist made five of the 10 judges think it was human. The judges did not have extensive computer expertise.

The contest was administered by the Cambridge Center for Behavioral Studies and hosted by the Computer Museum in Boston.

Applications for the 1992 competition must be postmarked by July 31. A maximum of 10 finalists will be selected by September 30 to participate in November's competition.

For information or applications, contact Robert Epstein, Director Emeritus and Program Officer, Loebner Prize Competition, Cambridge Center for Behavioral Studies, 11 Waterhouse St., Cambridge, MA 02138, e-mail psy9ain@buacca.bu.edu.

DARPA, Intel take aim at teraflops supercomputer

The US Defense Advanced Research Projects Agency and Intel have announced a coordinated effort to accelerate the development of high-performance computer systems capable of sustaining 1 trillion floating-point operations per second. The jointly funded effort, for which DARPA will contribute approximately $21 million over five years, directly supports the US High-Performance Computing and Communications Program.

Three critical technologies will form the basis for this teraflops initiative. Intel's advanced parallel software architectural design, developed as part of the company's Touchstone R&D program, will serve as the foundation for future generations of scalable parallel supercomputers. The fine-grained interprocessor communication and computing concepts implemented in Intel's iWarp real-time supercomputer systems will be enhanced and implemented in a new generation of submicron VLSI components. And, to provide the necessary advances in raw computational power, Intel will initiate the coordinated development of a new generation of high-performance floating-point RISC microprocessors derived from the Intel i860 series.

ADVERTISER INDEX

Alsys, Inc. 5

CACI Products Company 1

Computer Architecture Symposium '92 105

Design Automation Conference '92 91

GIGATEC, Inc. . . . . . . Cover III

Gordon Bell Prize 8

IEEE Computer Society Membership 87

IEEE Computer Society Ombudsman 64

IEEE Computer Society Press 16A-16B

Interactive Development Environments 9

PARLE '92 40

Quantum Software Systems Ltd. . . . . . . Cover II

SIGDOC '92 47

Software Engineering Institute 57, 61, 63, 65, 69

Systems Integration Conference 48-49

The University of Aizu 28

The University of Hong Kong 78

Wafer-Scale Integration Conference '93 96

Western Institute in Computer Science . . . . . . Cover IV

Classified Advertising 106-109

FOR DISPLAY ADVERTISING INFORMATION, CONTACT:

District Manager: D. Rodney Brooks; tel: (415) 905-0260; fax: (415) 896-1512.

Recruitment & Classified Advertising: D. Rodney Brooks; tel: (415) 905-0260; fax: (415) 896-1512.

Director of Sales: Randall L. Stickrod, South Park Media, 544 Second St., Suite 200, San Francisco, California 94107; tel: (415) 905-0260; fax: (415) 896-1512.

For production information, conference, and classified advertising, contact Heidi Rex or Marian Tibayan.

COMPUTER, 10662 Los Vaqueros Circle, PO Box 3014, Los Alamitos, California 90720-1264; tel: (714) 821-8380; fax (714) 821-4010.

PRODUCT INDEX

	RS#	PG#
Alsys, Inc.	2	5
Anadigics	121	92
Analogy	122	92
Artisoft	21	83
ATTO Technology	141	93
Bantam Books	22	82
Base Technology	23	83
CACI Products Company	—	1
CalComp Digitizer Products Grp.	43	90
Cirrus Logic	123	92
Cisco Systems	142	93
Colby Systems	37	89
Computer Dynamics	143	93
Cyclone Microsystems	144	93
Data Translation	145	93
Digimetrie	146	93
Digital Equipment	124	92
Dover Electronics	147	93
Force Computers	148, 149	93
GeoWorks	40	90
GIGATEC	4	C.III
Grass Roots	150	93
HDC Computer Corp.	28	85
Hitachi America	125	92
HumanCAD	29	86
IBM	38	89
IDE	3	9
Micro Touch Systems	36	89
MicroLogic Software	25	84
Microsoft Press	24	84
Motorola	126	92
National Semiconductor	127	92
Neural Computer	151	93
Newer Technology	128	92
Power Integrations	129	92
Quantum Software Systems Ltd.	1	C.II
RGB Spectrum	39	89
Shannon Associates	41	90
Slate Corp.	45	90
Stylos Development	42	90
Tektronix	35	89
Texas Instruments	130	92
Themis Computers	152	93
Traveling Software	44	90
Trident	131	92
Unison World Software	27	85
Weitek Corp.	26	85
White Technology	132	92
WICS	5	C.IV

COMPUTER

Reader Service Card

April 1992 issue (Card void after October 1992)
Print address below.

Name ____________________

Title ____________________

Company ____________________

Address ____________________

City ____________________

State/ZIP ____________________

Country __________ Phone __________

Reader interest
(Circle what you liked; add comments on the back)

- **A** Guest Editors' Introduction
- **B** Task-Flow Architecture
- **C** Design and Yield Evaluation
- **D** Configuring a 2D Array
- **E** Using Restructurable VLSI
- **F** Transducer Arrays
- **G** Prospects for WSI
- **H** Project Overviews
- **I** Letters to the Editor
- **J** CS News/Update
- **K** Product Reviews
- **L** New Products
- **M** ICs/Microsystems
- **N** Conferences
- **O** Call for Papers/Calendar
- **P** Book Reviews
- **Q** Open Channel

Product information
(Circle numbers for products and advertisements on which you want more information)

1	21	41	61	81	101	121	141	161
2	22	42	62	82	102	122	142	162
3	23	43	63	83	103	123	143	163
4	24	44	64	84	104	124	144	164
5	25	45	65	85	105	125	145	165
6	26	46	66	86	106	126	146	166
7	27	47	67	87	107	127	147	167
8	28	48	68	88	108	128	148	168
9	29	49	69	89	109	129	149	169
10	30	50	70	90	110	130	150	170
11	31	51	71	91	111	131	151	171
12	32	52	72	92	112	132	152	172
13	33	53	73	93	113	133	153	173
14	34	54	74	94	114	134	154	174
15	35	55	75	95	115	135	155	175
16	36	56	76	96	116	136	156	176
17	37	57	77	97	117	137	157	177
18	38	58	78	98	118	138	158	178
19	39	59	79	99	119	139	159	179
20	40	60	80	100	120	140	160	180

COMPUTER

Reader Service Card

April 1992 issue (Card void after October 1992)
Print address below.

Name ____________________

Title ____________________

Company ____________________

Address ____________________

City ____________________

State/ZIP ____________________

Country __________ Phone __________

Reader interest
(Circle what you liked; add comments on the back)

A Guest Editors' Introduction
B Task-Flow Architecture
C Design and Yield Evaluation
D Configuring a 2D Array
E Using Restructurable VLSI
F Transducer Arrays
G Prospects for WSI
H Project Overviews
I Letters to the Editor
J CS News/Update
K Product Reviews
L New Products
M ICs/Microsystems
N Conferences
O Call for Papers/Calendar
P Book Reviews
Q Open Channel

Product information
(Circle numbers for products and advertisements on which you want more information)

1	21	41	61	81	101	121	141	161
2	22	42	62	82	102	122	142	162
3	23	43	63	83	103	123	143	163
4	24	44	64	84	104	124	144	164
5	25	45	65	85	105	125	145	165
6	26	46	66	86	106	126	146	166
7	27	47	67	87	107	127	147	167
8	28	48	68	88	108	128	148	168
9	29	49	69	89	109	129	149	169
10	30	50	70	90	110	130	150	170
11	31	51	71	91	111	131	151	171
12	32	52	72	92	112	132	152	172
13	33	53	73	93	113	133	153	173
14	34	54	74	94	114	134	154	174
15	35	55	75	95	115	135	155	175
16	36	56	76	96	116	136	156	176
17	37	57	77	97	117	137	157	177
18	38	58	78	98	118	138	158	178
19	39	59	79	99	119	139	159	179
20	40	60	80	100	120	140	160	180

COMPUTER

Reader Service Card

April 1992 issue (Card void after October 1992)
Print address below.

Name ____________________

Title ____________________

Company ____________________

Address ____________________

City ____________________

State/ZIP ____________________

Country __________ Phone __________

Reader interest
(Circle what you liked; add comments on the back)

A Guest Editors' Introduction
B Task-Flow Architecture
C Design and Yield Evaluation
D Configuring a 2D Array
E Using Restructurable VLSI
F Transducer Arrays
G Prospects for WSI
H Project Overviews
I Letters to the Editor
J CS News/Update
K Product Reviews
L New Products
M ICs/Microsystems
N Conferences
O Call for Papers/Calendar
P Book Reviews
Q Open Channel

Product information
(Circle numbers for products and advertisements on which you want more information)

1	21	41	61	81	101	121	141	161
2	22	42	62	82	102	122	142	162
3	23	43	63	83	103	123	143	163
4	24	44	64	84	104	124	144	164
5	25	45	65	85	105	125	145	165
6	26	46	66	86	106	126	146	166
7	27	47	67	87	107	127	147	167
8	28	48	68	88	108	128	148	168
9	29	49	69	89	109	129	149	169
10	30	50	70	90	110	130	150	170
11	31	51	71	91	111	131	151	171
12	32	52	72	92	112	132	152	172
13	33	53	73	93	113	133	153	173
14	34	54	74	94	114	134	154	174
15	35	55	75	95	115	135	155	175
16	36	56	76	96	116	136	156	176
17	37	57	77	97	117	137	157	177
18	38	58	78	98	118	138	158	178
19	39	59	79	99	119	139	159	179
20	40	60	80	100	120	140	160	180

Editorial comments

I liked:

I disliked:

I would like to see:

For reader-service inquiries, see other side.

PO Box is for reader-service cards only.

PLACE
POSTAGE
HERE

COMPUTER
Reader Service Inquiries
PO Box 16508
North Hollywood, CA 91615-6508
USA

Editorial comments

I liked:

I disliked:

I would like to see:

For reader-service inquiries, see other side.

PO Box is for reader-service cards only.

PLACE
POSTAGE
HERE

COMPUTER
Reader Service Inquiries
PO Box 16508
North Hollywood, CA 91615-6508
USA

Editorial comments

I liked:

I disliked:

I would like to see:

For reader-service inquiries, see other side.

PO Box is for reader-service cards only.

PLACE
POSTAGE
HERE

COMPUTER
Reader Service Inquiries
PO Box 16508
North Hollywood, CA 91615-6508
USA

PRODUCT REVIEWS

Editor: Richard Eckhouse, University of Massachusetts at Boston
Send review submissions to J.M. Jagadeesh, College of Pharmacy, Ohio State University, Columbus, OH 43210

More about fantastic LANtastic

Richard Eckhouse, University of Massachusetts at Boston

As many of you know, I have been a LANtastic user for many years and have reviewed these products several times in *Computer*. Like many of you, I work in a multi-PC environment. Each PC is dedicated to a particular user or project, and each needs to share files and devices with at least one other PC. Years ago, I might have accepted a "sneaker net" to exchange floppies — but not anymore. I now expect to access every PC and every device at a cost that adds very little to the price of my original investment. Artisoft has continued to meet these needs with increasing functionality and decreasing price.

When I heard about Artisoft's new AE-3 Ethernet adapter card, version 4.0 of its Network Operating System (NOS), and the network management tool LANtastic for Windows, I immediately requested a review copy of each. As a Windows 3.0 user, I'd been unable to use my PC as a server under NOS Version 3.0. I've used Artisoft's constellation of products for months now — having worked through some minor problems — and I am very pleased with my ever-growing network. In addition, I recently designated one machine for archival storage by placing the Micropolis 1664 SCSI drive on it (see *Computer*, October 1991, pp. 85-86). That strategy has greatly reduced the storage demands on other machines and has made it considerably easier to maintain software so that one copy is current.

As I write this, I am running Ami Pro under Windows with Norton Desktop for Windows, with my 486/33 machine acting as a server. I have a simple batch file going on another machine that performs an endless loop asking for the root directory of my machine. LANtastic is operating transparently and unobtrusively, and — except for a few chirps coming from the other machine indicating it is waiting for network transfer — I can't tell what's happening behind the scenes. This is the way my LAN *should* work. Now I'll describe the pieces that make it happen that way.

AE-3 Ethernet Adapter

The AE-3 adapter is certified by Novell and works with thin/thick coaxial or twisted-pair Ethernet cabling by automatically switching between the two. It is compatible with the AE-2 (in my LAN I use both) and offers the same features (8/16-bit mode with 16 Kbytes of RAM upgradable to 64 Kbytes, NE2000 compatibility, and compliance with the IEEE 802.3 Ethernet standard). You can also use it on a hard-diskless PC with an optional ROM boot kit. The hardware includes the device driver and application-independent LANBIOS software, as well as the LANcheck diagnostic tool.

The user manual's thorough instructions range from explaining kit contents to setting board jumpers to attaching the network cabling. The appendixes include testing and error messages, and technical specifications.

The jumper setting includes selecting the IRQ, the I/O port address, and the Artisoft or NE2000-compatibility mode. You can also enable the remote-boot ROM, specify the cable length and type, and handle a nonstandard ISA bus. Two unusual features are the inclusion of IRQs 10 and 15 as well as 2 through 7, and a solution for the nonstandard bus problem. In the former case, I did have one system with all the lower IRQs already used; and in the latter case, I put an AE-3 in a Toshiba 5200 with its nonstandard bus. I was very pleased that Artisoft thought of these potential problems in advance so that I merely had to read the manual to solve them.

The low-level driver software, AEX, includes command-line switches for specifying the I/O port address, the IRQ, the multiplex interrupt number, the packet size, and the number of transmit buffers. A Remove option lets the user eliminate this TSR program. The Xerox option instructs the AEX software to send and receive data packets in the Xerox format, as well as to receive IEEE 802.3 data packets.

There are a great many command-line switches for AILANBIOS, the adapter-independent LAN BIOS software. You can specify time-outs, network control blocks, number of network sessions, and number of networks. This software also includes a Remove command. Both AILANBIOS and AEX command-line switches can be stored in an indirect file rather than as part of a command. A Verbose option that provides detailed information about the AILANBIOS configuration was a valuable addition to my network-startup batch file.

You can install multiple network cards in the same machine (up to four AE-3s or -2s), but the LANtastic network software does not currently support bridging. In years of use, I've never had an adapter fail. My only problems have been with other boards that do not have the flexibility of the AE-3 (or -2) for changing the I/O port address or the IRQ. For $349 (or less through mail order), this is a very inexpensive way to bring a new machine on line with an existing network. But you can also buy a bundled package consisting of two AE-3s, AILANBIOS, NOS 4.0, plus 25 feet of coaxial cable for $799. I rate this as a best buy.

LANtastic NOS Version 4.0

As indicated earlier, version 3.0 of NOS presented a problem: It would not let machines act as servers while running Windows in enhanced mode. While version 4.0 corrected that problem, a new problem occurred in conjunction with using Norton Desktop for Windows (NDW). That problem was resolved only recently through an update from NDW's company, Symantec. You will probably encounter some minor annoyances, nevertheless. For instance, when I set LASTDRIVE = L, my networked drive (called F) came up as a drive F icon under NDW; but when I clicked on the icon, drive F magically became drive H as the root of the tree and no further file operations were possible in NDW. Symantec tech support suggested setting the LASTDRIVE to Z, and that fixed the problem.

Once this was fixed, I really saw the value of running NDW with LANtastic. Moving files between machines was as simple as opening two drive windows and selecting the files you want from the other machine in the drive window of the one you are using. This process becomes so convenient and fast that you really aren't aware your machine is running within a network environment.

The significant changes between versions 4.0 and 3.0 (reviewed in *Computer*, December 1990, p. 88) — besides full support for Windows 3.0 in all modes — include

- an improved installation procedure (already very well done, in my book);
- the capability to unload NOS from memory;
- UPS support;
- an Alone program to dedicate a machine as a server, sidestepping DOS;
- a fast network-caching program that caches both reads and writes (replacing Windows' SmartDrive);
- remote processing (to use NOS to run a program on a server);
- improved group accounts; and
- faster file transfers.

Given the low cost of an upgrade from version 3.0 ($50 for the LANtastic Z, Ethernet, and 2-Mbyte versions, and $99 for the adapter-independent version), these benefits are very reasonably priced.

For those of you unfamiliar with NOS software, it offers peer-to-peer networking with a minimal requirement for DOS memory (12 to 50 Kbytes, depending upon configuration). Less memory space is required if your machine is only a workstation that accesses files and devices on other machines; more is needed if your machine acts as a server. Either way, you get a very powerful network environment that supports

- full compatibility with most NETBIOS adapters,
- access control of disk drives and subdirectories by user or group,
- audit trails,
- e-mail with a chat facility,
- up to 300 concurrent users per service,
- remote control and configuration of machine functions,
- mainframe-like security,
- file and record locking, and
- CD-ROM support.

These features aren't new in version 4.0 but are improved with each release.

Automated installation and a set of manuals offer a quick way to set up the hardware and get the network running. Two utilities, Net and NetMgr, provide necessary tools for the network user and the system administrator. The Net program provides for login/logout, network-disk-drive and printer access, printer queue management, mail services, the chat facility, user-account management, and a display of server activity. The NetMgr program provides for managing individual/group accounts and shared resources; setting server start-up parameters; and maintaining audit trails, queues, passwords, boot images, and control directories. A simplified version of Net, called LANPUP (LAN pop-up utility program), comes as a TSR.

This excellent network system includes all the hardware and software needed to set up a full-featured LAN. It represents a significant value for the investment and should be very seriously considered by anyone thinking of building a LAN. I have recommended it many times, and every subsequent user has thanked me for this advice.

Additional reading

As Artisoft's LANtastic becomes more widely used and richer in features and functionality, it seems quite natural that someone would write a book entitled *Running LANtastic.* This 624-page 1991 Bantam Computer Book by Adrian J. King fully covers most of Artisoft's network products and, at $27.95, is modestly priced.

You might ask, "What can I get out of this book that I can't find in the product manuals?" Lots, because King has managed to cover the material *and* add value through explanations of why and how things work. And for the particularly technical reader, he includes a chapter on the LANtastic programming interface, information not normally found in the materials Artisoft packs with its products.

I must admit that I approached this book with a sort of ho-hum attitude. I have been a LANtastic user for many years and have kept up with the many new releases of software and hardware. When I read the opening chapters that cover the basics of PC networking as viewed from the LANtastic perspective, I didn't gain much insight. King anticipates this in his introduction. But when I delved into the following chapters having to do with the details of installing, configuring, and using a LANtastic network, I became much more interested and quickly picked up some new ideas. Yet another section deals with more advanced features of LANtastic, including performance tuning. The book also covers some products I haven't used recently, LANtastic Z and Sounding Board. It also includes three appendixes about the Net command reference, the programming interface, and a short bibliography.

King's writing style allows ideas to flow naturally and be rapidly assimilated. The book includes a number of figures and is punctuated with notes and chapter summaries. King was an assistant general manager of the Network Business Unit of Microsoft and currently is vice president of engineering at Artisoft. His background illuminates this technically complete book. At the same time, it is a pleasure to read and offers a wealth of knowledge not to be found elsewhere. I have added this book to my treasured list of technical references.

Readers can contact Bantam Books at 666 First Ave., New York, NY 10103; (212) 765-6500; or circle Reader Service Number 21.

— Richard Eckhouse, University of Massachusetts at Boston

LANtastic for Windows

At first, I mistakenly thought LANtastic for Windows (LFW) was a replacement for LANtastic NOS. Rather, it is a graphical implementation of the network utilities that come with NOS. Thus, you must own NOS and Windows 3.0, and have DOS 3.1 or later to support the software.

Even if all that LFW did was to make the Net and NetMgr programs run as Windows applications, I would still pay the $299 per-network price. First, you can login/logout, select/connect drives or printers, chat or send mail to another LANtastic user, and examine queues without having to open a DOS window. Second, you can click on your selections (or click and drag), so you don't have to remember Net commands or the network names for things. Third, a graphical display of server activity is more pleasing than the raw numbers themselves. But last — and most important — is that Artisoft has made LFW dynamic data exchange (DDE)-intelligent.

LFW includes DDE across the network. Though barely mentioning it in the user manual, Artisoft does devote Appendix A to using LANtastic with other Windows applications. That DDE capability takes on a new dimension when a program on one machine can exchange messages and data with a user on another machine. To quote the manual:

> You can use Word for Windows or even Windows Paintbrush to create mail messages. When the remote user receives the message, all he or she will have to do is double-click on the mail item in the Queues window and the file will open in the application in which it was originally created.

To provide this capability, LFW supports five DDE commands: CreateWindow, MailFile, PrintFile, SelectServer, and ExitNet. Also provided are macros for Word for Windows and Excel, along with a description of how to use them.

A side benefit of using LFW is that it maintains its own INI file. This means that you can tell the system to automatically save the latest configuration in the Windows environment. It also means that the batch file for network startup can be abbreviated to merely load the hardware drivers and NOS software, leaving the full configuration up to the Windows user. Log-in becomes a simple click of the mouse, and network devices are added or deleted by using the drag-and-drop procedure.

Although LFW is feature rich, the learning curve is simple because this software mimics what was previously available in its DOS equivalent. Installation is automatic, with LFW modifying the Windows INI files to ensure that LFW is loaded each time your network software and Windows are started up. The manual, while brief, is thorough and nicely indexed.

Summary

Artisoft's LANtastic continues to be fantastic. For a modest price, I can conveniently access all the resources connected to my LAN. Network commands are intuitive, with on-line help whenever I need it. You get both the hardware and software to implement a feature-rich LAN that has more than you would normally expect. From my perspective, the savings and flexibility are so compelling that it makes sense to buy LANtastic even if you are only planning on a two-node network.

Readers can contact Artisoft at Artisoft Plaza, 575 E. River Rd., Tucson, AZ 85704; (602) 293-6363; fax (602) 293-8065; or circle Reader Service Number 22.

Review notes

Liana. This object-oriented C++-like language allows you to create Microsoft Windows applications. However, the differences between Liana and C++, which are clearly documented in two pages at the end of the manual, do exist. In general, I found these differences to be in Liana's favor. For example, Liana's switch expressions are not limited to integers, and all arrays are dynamic. However, you might run into problems converting an existing C++ program; for instance, union, enum, and function prototypes are not supported in Liana.

Liana does not use anything like the Microsoft Windows Software Development Toolkit (SDK). The 421-page manual includes documentation for the language structure, all functions, and built-in classes. It also explains how to interface with existing document-description languages (although you cannot write a DDL with Liana) and provides 62 examples. The examples cover everything from the construction of dialogue boxes and menus to DDE clients/servers and using the PC speaker. The distribution disk contains the source code for the examples.

Writing a Liana program was very simple. It was easy to create dialogues and menus, and interface my program with them. The only limitation I saw was that all dialogue boxes and menus are created dynamically. It takes a long time to create and display the complex dialogue box. Neither can Liana use a standard Windows resource, which means that dialogues and menus cannot be created before execution.

Both the compiler and linker are DOS programs. Although this was disappointing, it wasn't a big problem. Because a Liana program can be written using any ASCII text editor, I used Windows Notepad to write my programs. I had DOS running in a small window to support the Liana compiler and linker, and my file manager in another window to launch my program. (Users with a 286-based system will find the DOS compiler/linker more problematic because the DOS window must be full screen.)

I had three difficulties with Liana. First, after about eight invocations of the compiler/linker program, I kept getting an error message that the compiler/linker had violated system integrity. The DOS window then closed. I could open a new DOS window and proceed — but this process was very strange. Second, Liana does not include any kind of debugger. The manual does discuss debugging, but it all boiled down to inserting output statements in the code. I cannot see doing any serious programming without some kind of debugger. Third, Liana garbage collection is not very complete. The result was that Liana programs slowly acquired all my system's free memory. After a while (five minutes in some cases), the lack of free memory caused Windows to start behaving very strangely, and I had to reboot my system.

Liana holds great promise for someone who wants to program in C++ but doesn't want to deal with the massive documentation and very steep learning curve of Microsoft Windows' SDK or Borland's C++. However, the $495 program needs a debugger and improved memory management before it can be considered a serious alternative.

Readers can contact Base Technology at 1543 Pine St., Boulder, CO 80302; (303) 440-4558; or circle Reader Service Number 23.

— Noah Davids, Stratus Computer

All in one place. Finding information and technical details about the IBM PC, MS-DOS, and other products is often tedious because the facts are spread over many volumes. Therefore, when you are programming or debugging at the system level, you normally have to consult a number of references for needed information about BIOS functions or the differences between DOS versions — or the pinouts of some device controllers and the interpretation of bits found in their registers. However, you can get all the information you need quickly from the *Programmer's PC Sourcebook* by Thom Hogan from Microsoft Press, La Vergne, Tenn., 1991, 808 pp.

This collection of basic hardware and software information contains about 50 percent more material than its first edition does. New topics include MS-DOS 5.0, Windows 3.0, CD-ROM Extensions, MS Mouse Driver 8.0, LIM 4.0, the new EMS capabilities, EISA, Hayes Modem commands, and the Virtual Control Program Interface. This edition also details the PS/2 Models 80, 90, and 95; the 386, 387, and 486; and the Weitek 3167 and 4167.

The first part of the book provides some general information for programmers. Tables for numeric conversions and common data formats are well organized and easily retrievable. The book also presents extensive and easily understood lists and tables on DOS commands, utilities and function calls, BIOS and DOS extension calls, other interrupts, and Windows (including input devices, files, formats, resource script files, the contents of the WIN.INI file, functions, and utilities). General information about PC hardware, such as motherboards, keyboards, video adapters, peripherals, chips, connectors, buses, and pinouts, are all included. Information about older systems has been retained from the first edition, and equivalent features of earlier software versions are identified.

I had frequently consulted the first edition, but it went out of date as new hardware and software became available. Consequently, I stopped using it years ago. Having the second edition makes me feel much more comfortable about finding answers to my numerous questions as I write code for the PC. At $39.95, the book is an excellent and extremely helpful reference. It boggles my mind to think of how many manuals and references would clutter my computer table without the *Programmer's PC Sourcebook* at my side.

Readers can contact Microsoft Press at PO Box 7005, La Vergne, TN 37086; (800) 677-7377; or circle Reader Service Number 24.

— Qing Zhuang, MOCO Inc.

Display Faces. I hate to admit it, but I have become a font junkie. It started long ago when I first reviewed MoreFonts (MF) using the Kyocera F-1000A, a plain vanilla laser. It hasn't gotten any better now that I use the QMS-PS 410, a PostScript laser. I just like my correspondence to look interesting. With the large number of add-on typefaces, it's hard to resist the temptation to play with fonts. I could blame MicroLogic Software for sending me a copy of Display Faces, but the truth is, I asked for it!

If you go back to the December 1991 issue of *Computer* (p. 89), you will find my latest review of MoreFonts, a type-scaling system that works with WordPerfect products, Microsoft Word, PageMaker, Ami Professional, and Windows. I enjoyed reviewing that program with its 28 scalable typefaces and use it all the time. But I wanted more, particularly the Mechanical typeface found in Display Faces.

The package arrived and, like all of MicroLogic's software, it was easily installed by simply typing "A:Hello" at the prompt. However, I needed to add my new typeface to MoreFonts, so I clicked on the MF icon, clicked on the new typeface I wanted, closed MF, and restarted Windows. Things couldn't be simpler. I didn't even need to look at the manual (which is complete, as all the MF manuals are). Within minutes, I had the typeface I wanted and was developing a new form using Ami Pro.

Interestingly, MicroLogic claims Display Faces not only works with MoreFonts but also with many other type-scaling programs including the Agfa/Intellifont, Atech, GeoWorks Ensemble, and PostScript Type-1 formats. A utility comes with the package to convert to any of these formats, which adds a lot to this $149.95 package (street price is about $100).

I remain impressed with MicroLogic Software and the innovative, value-packed products it releases.

Readers can contact MicroLogic Software at 1351 Ocean Ave., Emeryville, CA 94608; (510) 652-5464; fax (510) 652-7079; or circle Reader Service Number 25.

— Richard Eckhouse, University of Massachusetts at Boston

Power for Windows. I've always thought of Weitek as the manufacturer of accelerator chips for the IBM PC and compatibles. So when I heard Weitek was making a user-interface controller chip to speed up the Windows GDI (graphics display interface) functions, I wasn't surprised. But seeing the company enter the VGA board market did catch my interest, so I stopped by its booth at Comdex to find out more about Weitek's Power for Windows.

Weitek's board is a 16-bit VGA graphics adapter based on its W5086 chip, whose line-drawing and bit-block transfer (bitblt)-acceleration features should increase Windows 3.0 performance by 25 to 200 percent. It offers resolutions of up to 1,024 × 768 in 16 colors. The half-size card contains a fixed amount of memory (512 Kbytes) and only three jumpers that may be changed: ROM and RAM 8- and 16-bit modes, and the use of IRQ9 (although most of us won't change these). It also uses two non-standard VGA register ports addressed at 30Ah and 32Ah, which caused a minor conflict with my Artisoft Ethernet adapter board at 300h because the Weitek user manual doesn't make it clear that the non-standard addresses actually reach down to 300h.

Weitek's friendly tech support identified the problem immediately. After installing the board, users are instructed to run the Modeset utility that sets both the board (to match the characteristics of the monitor being used) and the video mode. I set the video to 132 × 50 mono mode and was pleased to see that the directory command scrolled down fully 50 lines. The display is slower than in normal 80 × 25 color mode but not nearly so slow as when I put it in 1,024 × 768 × 16-color graphics mode. Obviously, you pay for what you get.

Installation of the Windows software is automatic and the most intelligent I have seen (and I have tested a great many boards). The installation program modifies the SETUP.INF file, so that instead of replacing it, you can preserve all other display adapter selections you had previously established. I hope that other display adapter manufacturers adopt this idea because having multiple SETUP.INF files on floppies is so inconvenient.

After installing the Weitek software by running the Windows setup, I next brought up Windows in 1,024 × 768 resolution. My typical applications are Ami Pro and Excel; my first impression was that the display scrolling in both applications was considerably faster than it had been. One minor inconvenience was that the Ami Pro background flickered. But a change of my Windows color scheme easily fixed that. And while I had been running Windows in 256-color mode, it didn't make any difference to me that Weitek supports only 16.

Using the same set of benchmarks I used to measure Panacea's WinSpeed software accelerator (see *Computer*, March 1992, pp. 91-92), I confirmed

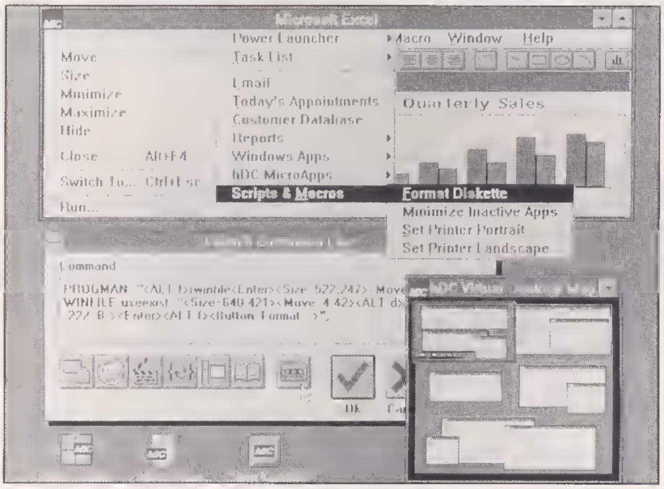

Power Launcher pop-up menu.

my subjective opinion. Compared to an original STB Ergo board running in 1,024 × 768 × 16 resolution, the Weitek board turned in a performance that ranged from 22 percent slower to 3,367 percent faster. The "draw arcs" was the slowest test, and the "screen bitblt source invert" was the fastest. Panacea's software sped up the STB board so that the differences covered a smaller range from 66 percent slower to 1,085 percent faster.

The only complaint I have with Weitek's board is that it's an I/O port hog. I used my machine's address space from 300h to 330h and forced myself to carefully check whether the other devices could be relocated, because the Weitek board itself can not. Aside from that, how would I rate this $299 board? When asked by a good friend if I would give it to him after finishing my review, I replied, "No way; it stays in my machine!"

Readers can contact Weitek Corp. at 1060 East Arques Ave., Sunnyvale, CA 94086; (408) 738-8400; or circle Reader Service Number 26.

— Richard Eckhouse, University of Massachusetts at Boston

Avagio 2.0. As a new PC user, I jumped at the opportunity to try out the Avagio 2.0 desktop publishing program. I wanted to create newsletters or other publishing work instead of relying on a copy center.

Installation is fairly easy; you type "install" and follow the instructions on the screen. What I like about Avagio is that it lets me create different designs by using the large collection of clip art in the package (like balloons, handshake, space shuttle, or airplane, just to name a few). I can choose size, black/white, and a variety of mix-and-match colors. The graphics can be turned in any direction and placed anywhere on the screen. They can be enlarged or reduced, while all other elements on the screen remain untouched.

You can use the Mingle feature to create special effects for text and graphics. It lets you see how they would look when one is placed over the other. Color blending also comes with this feature.

The greatest part of using this program is that you don't need much publishing talent. Plain and simple, this system is easy to use. All you have to do is follow the manual or screen instructions. After a very short time, you can become a confident user.

Avagio costs only $149.95 and runs on any IBM or compatible 286, 386, or 486, as well as on the Tandy 1000.

Readers can contact Unison World Software at 1321 Harbor Bay Pkwy., Alameda, CA 94501; (510) 748-6670; fax (510) 748-6964; or circle Reader Service Number 27.

— Marie M. Daniel, University of Massachusetts at Boston

Power Launcher. We could probably all agree that any windowing system is better than none. But let's face it, Microsoft Windows 3.0 needs help. The HDC Power Launcher offers it in a collection of tools.

Window-manager software often suffers from a designer's belief that it is running on a small system. The software may operate on lurking assumptions that work well only on small screens or with small file systems with a limited amount of memory (so that you don't have very many applications launched at one time). Windows is no exception. The HDC Windows enhancement package provides ways for you to minimize wandering about in your directory trees and reduce the amount of "mousing around" necessary for launching applications. Its most useful aid is controlled by a button provided on the title bar of every application you launch. This permits easy switching to any other task you have running (including the Program Manager, which you will use rarely once you start configuring your own launch scripts with this software).

Power Launcher and its associated script builder enable you to specify a complex sequence of applications and their parameters. According to the documentation, under Windows 3.1, these applications can exchange data with each other via the 3.1 dynamic data-exchange mechanisms even before you see a window open in the final application in the chain. (Unix users are familiar with executing pipelines of applications.) I tested Power Launcher only under Windows 3.0, but the pipelining support is enough to make me want to use it under 3.1!

Power Launcher includes three other potentially useful applications. The first two are for remapping keys and mouse buttons (globally or associated with a particular application). The other is for building button-based user interfaces that can be placed at any convenient position on your desktop (including anywhere on Power Launcher's large virtual desktop).

The keyboard mapper disappointed me: I had hoped to unify the key bindings for my Windows-based editors so they would be the same as I use in the emacs editor. I used the macro binding facility to map CTRL-f to the same effect as the right-arrow key. The result was to bind to the correct "forward character" in Notepad but "forward word" in Windows Write. Worse, the application cannot sense the state of the control key: You must release and press it again for each application of the command. HDC's prompt and helpful technical support was unable to reproduce the misbehavior in the key binding. It may be an interaction between HDC and my Gateway keyboard. The company did state that the control key sensing will be changed in the next release (scheduled for Spring 1992 Comdex) to coincide with more typical application behavior.

Despite the disappointing keyboard binder, I feel this product's launching utility is well worth the street price of about $65 (list price $99.95.)

Readers can contact HDC Computer Corp. at 6742 185th Ave. NE, Redmond, WA 98052; or circle Reader Service Number 28.

— Robert Morris, University of Massachusetts at Boston

Mannequin. As frequent CAD users and most computer illustrators know, it's difficult and time-consuming to create a realistic graphic image of any object — never mind a true-to-life, functional human figure. And in some areas of business, research and development, and manufacturing, ergonomic correctness rivals political correctness as a standard. For those who need to create systems and devices with human factors in mind — or to create an anatomically and functionally sound human figure for presentation graphics and animation — Mannequin from HumanCAD can ease the laborious process.

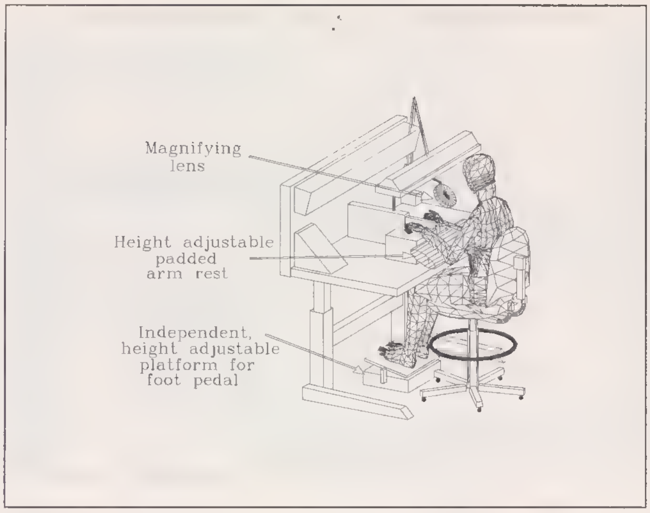
Magnifying lens
Height adjustable padded arm rest
Independent, height adjustable platform for foot pedal

Mannequin drawing.

The Mannequin software package quickly creates ergonomically correct human figures, or "mannequins," that can be selected by nationality, age, gender, body size, and statistical category. The mannequins can be set in a number of poses (with hand gestures) and edited to hold a desired pose, or positioned with other graphical objects. Keeping with the company's claims about ergonomical correctness, Mannequin makes only human-like motions and can be animated frame by frame with the program-recording features or can walk along a specified path. Users can also reach for a selected target for a quick animation display to perform an easy test of human fit.

Aiding human factors considerations is a vision function that displays through a grid the optimal and/or maximal range of the mannequin's vision. Although this is effective, it is enhanced with Human View, an option that lets the user see the drawn scene through the mannequin's eyes. Another visual guide is the Range of Motion selection, which displays a grid of the maximum possible range of motion for the right or left hand or foot, in horizontal and/or vertical planes.

An impressive torque calculator displays the biomechanical torque on major joints and compares torques of any two mannequins. Finally, if the user wishes to view the anthropometric data of mannequins, the software has a printable display of body measurements and range of motion, in numbers, of all body parts.

The users who would probably benefit most from this program are the artists or engineers who are already established in another compatible graphics/CAD/animation program. These users could import the mannequins to their own drawings for human-fit evaluation or for inclusion in a presentation-graphics or animation scene. After all, Mannequin was originally intended to supplement other graphics programs.

Mannequin, however, can function on its own. The program has 2D and 3D drafting capabilities that are quite impressive for a piece of software not intended as a pure CAD program. However, the editing functions should be reconsidered. A much-needed addition is an Undo function, or for that matter, anything that lets users edit mistakes without having to erase entire objects and start over. This is particularly true of the 2D and 3D objects but also applies to the mannequin form, which is difficult to reposition once set in a pose.

The grouping and selection-by-shape commands (3D objects, mannequin, text, etc.) seemed next to useless. Often I had to select a function twice before it would operate. My guess is that a conflict exists between "Press right button to select" and "Press right button to exit." I suggest that users remain with a graphics program they are comfortable with and try exporting their designs to Mannequin.

Although a capable program, Mannequin is not a polished one. Mannequins cannot be smoothed from their wire-frame form, and the closest the software can come to rendering is a shading function, which shades according to a user-placed light source. Animation, although more than functional, can be choppy and would best be cleaned up by using another graphics package. Users also have to look to other software (specifically Autodesk 3D Studio) for such features as smoothing the mannequin for a more streamlined look, animating the motions between two selected poses, distorting mannequins, and creating faces — all features that I would like to see within the package. I was disappointed that this expensive, memory-intensive software could not function as an effective closed system.

The company recommends using an expanded memory manager (such as QEMM 5.1) to allow for the software's 2-Mbyte-plus expanded memory. I used EMM386, resident to MS-DOS 5.0, with extra memory allocated; but I experienced a number of program crashes while experimenting to determine the amount of necessary memory. Mannequin is allegedly Windows 3.0 compatible; but after many struggles with the included .PIF file, File Manager, and elusive installation instructions, I found otherwise.

Actually, the instructions need the most work. Lacking a tutorial, the documentation is sometimes vague, often describing what can be done rather than how to do it. The user guide is out of date, and you have to turn to the present software addendum to have some quirks explained. On-line help is a little better, but in some areas it's just an echo of the manual.

Do I recommend the package? It could be a good choice if you already have an optimal graphics system yet would like help with human drawings. It's the only software of its kind for the PC. It's also a powerful program with sound statistical and ergonometric information and a lot of brilliant ideas behind it and . . . potential. While using the software, I constantly got a feeling of "It's not done yet." At $499, unless cheap upgrades appear in the very near future, I'd wait for the next version.

Readers can contact HumanCAD at 1800 Walt Whitman Rd., Melville, NY 11747; (516) 752-3568; or circle Reader Service Number 29.

— Karin Feeney, MOCO Inc.

IEEE COMPUTER SOCIETY

1992 Membership/Subscription Application

THE INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, INC.

BENEFITS

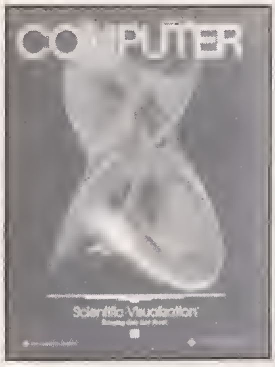

Receive *Computer* automatically with membership. Written and refereed by experts, it features articles on the latest developments in computer technology, applications, and research, as well as survey and tutorial articles on topics covering the entire computer field.

Computer's popular departments include New Products, Product Reviews, Book Reviews, Standards, Open Channel (a reader forum), and Computer Society News.

Other membership benefits include the following:

- Discounts of up to 50% on **Computer Society Press books and videos.** Over 700 titles covering a broad spectrum of computer science topics including networking, communications, advanced systems, image processing, security, artificial intelligence, and visualization
- Discounts on registration to over 100 conferences, workshops, and tutorials each year
- Low subscription rates on special-interest periodicals
- Opportunity to participate in any of our 32 technical committees—networks of professionals with common interests in computer hardware, software, and applications
- Opportunity to participate in the development of more than 200 standards projects sponsored by the society
- Low member rates on COMPMAIL, the electronic mail service especially designed for computer professionals

Schedule of Fees

To join: see item 1, 2, or 3.
To subscribe: see item 4.

Membership dues and periodical subscriptions are annualized to, and expire on, December 31. Pay full- or half-year rate depending on date of receipt by the Computer Society as indicated below.

PRICES EXPIRE 12/31/92

	Half Year Mar 1-Aug 31	Full Year Sept 1-Feb 28
1 I don't belong to the IEEE and I want to join just the Computer Society	☐ $27	☐ $ 54
2 I want to join both the Computer Society and the IEEE*		
I reside in Region 1-6 (United States)	☐ $58.50	☐ $117
I reside in Region 7 (Canada)	☐ $53.50	☐ $107
I reside in Region 8 (Europe, Africa, or the Middle East)	☐ $53.00	☐ $106
I reside in Region 9 (Latin America)	☐ $49.50	☐ $ 99
I reside in Region 10 (Asia and Pacific)	☐ $48.50	☐ $ 97

*ACM members who join both IEEE and the Computer Society may deduct $5 off the full-year rate; $2.50 off the half-year rate.

	Half Year	Full Year
3 I already belong to the IEEE and I want to join the Computer Society	☐ $11.00	☐ $22

IEEE Member Number ______________

4 SPECIAL-INTEREST PERIODICALS for new or current members

	issues per year	Half Year	Full Year
IEEE Computer Graphics and Applications	6	☐ $12.50	☐ $ 25
IEEE Design and Test	4	☐ $11.00	☐ $ 22
IEEE Expert	6	☐ $10.00	☐ $ 20
IEEE Micro	6	☐ $11.50	☐ $ 23
IEEE Software	6	☐ $12.50	☐ $ 25
Transactions on:			
Computers	12	☐ $12.00	☐ $ 24
Knowledge and Data Engineering	6	☐ $ 9.50	☐ $ 19
Parallel and Distributed Systems	6	☐ $ 9.50	☐ $ 19
Pattern Analysis and Machine Intelligence	12	☐ $12.00	☐ $ 24
Software Engineering	12	☐ $11.00	☐ $ 22
IEEE Annals of the History of Computing	4	☐ $ 8.00	☐ $ 16

☐ Visa ☐ MasterCard ☐ American Express

☐☐☐☐☐☐☐☐☐☐☐☐☐☐☐☐

Charge Card Number (Minimum Charge $10)

Mo. Yr.
☐☐☐☐
Exp. Date

Checks drawn in local currency on a bank in the country of origin accepted from members in Australia, Austria, Belgium, Canada, Denmark, France, Germany, Ireland, Italy, Luxembourg, Netherlands, Norway, Portugal, Spain, Sweden, Switzerland, UK, USA, and Yugoslavia.

Periodicals total $________

CA, DC residents add applicable sales tax $________

Membership fees $________

Total $________

CANADIAN residents add 7% GST; BELGIAN residents add 6% VAT $________

Total enclosed $________

I hereby make application for Computer Society membership and, if elected, will be governed by IEEE's and the society's constitutions, bylaws, and statements of policies and procedures.

MAILING ADDRESS

Full signature ________ Date ________

☐ Male ☐ Female

First name ________ Middle initial(s) ________ Last name ________ Date of birth ________

Street address ________ City ________ State/Country ________ Zip ________

EDUCATION (highest level completed)

Course ________ Degrees received ________ Date ________

Name of educational institution ________

OCCUPATION

Title or Position ________

Firm name ________

Firm address ________

City ________ State/Country ________ Zip ________

ENDORSER (an IEEE member, Senior Member, or Fellow)

Endorser's signature ________

Name (print in full) ________ IEEE Member No. ________

Street address ________

City ________ State/Country ________ Zip ________

Mail or fax to appropriate Computer Society office:

EUROPE: 13, Avenue de l'Aquilon, B-1200, Brussels, BELGIUM. Fax: 32-2-770-8505
PACIFIC RIM: Ooshima Building, 2-19-1 Minami-Aoyama, Minato-ku, Tokyo 107, JAPAN. Fax: 81-3-3408-3553
ALL OTHERS: 10662 Los Vaqueros Circle, PO Box 3014, Los Alamitos, CA 90720–1264 USA. Fax: 714-821-4010

PC0492

IEEE COMPUTER SOCIETY®

A member society of the
Institute of Electrical and Electronics Engineers, Inc.

THE FOLLOWING INFORMATION IS AVAILABLE:

Contact the Publications Office;
to facilitate handling, please request by number.

- Membership application, student #203, others #202
- Periodicals subscription form for individuals #200
- Periodicals subscription form for organizations #199
- Publications catalog #201
- Compmail electronic mail brochure #194
- Technical committee list/application #197
- Chapters lists, start-up procedures #193
- Student scholarship information #192
- Volunteer leaders/staff directory #196
- IEEE senior member grade application #204
  (requires ten years practice and significant performance in five of those ten)

To check membership status or report a change of address, call the IEEE toll-free number, 1-800-678-4333. Direct all other Computer Society related questions to the Publications Office.

PURPOSE

The IEEE Computer Society advances the theory and practice of computer science and engineering, promotes the exchange of technical information among 100,000 members worldwide, and provides a wide range of services to members and nonmembers.

MEMBERSHIP

Members receive the acclaimed monthly magazine *Computer*, discounts, and opportunities to serve (all activities are led by volunteer members). Membership is open to all IEEE members, affiliate society members, and others interested in the computer field.

PUBLICATIONS AND ACTIVITIES

Computer. An authoritative, easy-to-read magazine containing tutorial and in-depth articles on topics across the computer field, plus news, conferences, calendar, interviews, and product reviews.

Periodicals. The society publishes seven magazines and five research transactions. Refer to membership application or request information as noted above.

Conference Proceedings, Tutorial Texts, Standard Documents. The Computer Society Press publishes more than 100 titles every year.

Standards Working Groups. Over 100 of these groups produce IEEE standards used throughout the industrial world.

Technical Committees. More than 30 TCs publish newsletters, provide interaction with peers in specialty areas, and directly influence standards, conferences, and education.

Conferences/Education. The society holds about 100 conferences each year and sponsors many educational activities, including computing science accreditation.

Chapters. Regular and student chapters worldwide provide the opportunity to interact with colleagues, hear technical experts, and serve the local professional community.

OMBUDSMAN

Members experiencing problems — magazine delivery, membership status, or unresolved complaints — may write to the ombudsman at the Publications Office.

EXECUTIVE COMMITTEE

President: Bruce D. Shriver*
17 Bethea Drive
Ossining, NY 10562
Phone: (914) 762-3251
Fax: (914) 941-9181

President-Elect: James H. Aylor*
Past President: Duncan H. Lawrie*

VP, Conferences and Tutorials: Barry W. Johnson (1st VP)*
VP, Educational Activities: Raymond E. Miller (2nd VP)*
VP, Membership Activities: Fiorenza C. Albert-Howard*
VP, Press Activities: Yale N. Patt*
VP, Publications: Harold S. Stone*
VP, Standards Activities: Gary Robinson†
VP, Technical Activities: Joseph Boykin†

Secretary: Ronald G. Hoelzeman*
Treasurer: Laurel V. Kaleda†
IEEE Division V Director: Bill D. Carroll†
IEEE Division VIII Director: Helen M. Wood†
Executive Director: T. Michael Elliott†

*voting member of the Board of Governors
†nonvoting member of the Board of Governors

BOARD OF GOVERNORS

Term Expiring 1992:
Alicja I. Ellis, Ronald G. Hoelzeman, Tadao Ichikawa,
C.V. Ramamoorthy, Sallie V. Sheppard,
Harold Stone, Akihiko Yamada

Term Expiring 1993:
Fiorenza Albert-Howard, Jon T. Butler, Michael C. Mulder,
Yale N. Patt, Anneliese von Mayrhauser,
Benjamin W. Wah, Ronald Waxman

Term Expiring 1994:
Mario R. Barbacci, Luis-Felipe Cabrera, Wolfgang K. Giloi,
Guylaine M. Pollock, John P. Riganati, Ronald D. Williams,
Thomas W. Williams

Next Board Meeting
June 5, 1992, 8:30 a.m.
Sheraton New Orleans, New Orleans, LA

SENIOR STAFF

Executive Director: T. Michael Elliott
Publisher: H. True Seaborn
Director, Conferences and Tutorials: Anne Marie Kelly
Director, Finance and Information Services: Tod S. Heisler
Director, Board and Administrative Services: Violet S. Doan
Assistant to the Executive Director: Sandra K. Pfau

COMPUTER SOCIETY OFFICES

Headquarters Office
1730 Massachusetts Ave. NW
Washington, DC 20036-1903
Phone: (202) 371-0101
Fax: (202) 728-9614

Publications Office
10662 Los Vaqueros Cir.
PO Box 3014
Los Alamitos, CA 90720-1264
Membership and General Information:
(714) 821-8380
Publication Orders: (800) 272-6657
Fax: (714) 821-4010

European Office
13, Ave. de L'Aquilon
B-1200 Brussels, Belgium
Phone: 32 (2) 770-21-98
Fax: 32 (2) 770-85-05

Asia/Pacific Office
Ooshima Building
2-19-1 Minami-Aoyama, Minato-ku
Tokyo 107, Japan
Phone: 81 (3) 3408-3118
Fax: 81 (3) 3408-3553

IEEE OFFICERS

President: Merrill W. Buckley, Jr.
President Elect: Martha Sloan
Past President: Eric E. Sumner
Secretary: Karsten E. Drangeid
Treasurer: Theodore W. Hissey, Jr

VP, Educational Activities: Edward A. Parrish
VP, Professional Activities: Arvid G. Larson
VP, Publication Activities: James T. Cain
VP, Regional Activities: Luis T. Gandia
VP, Standards Activities: Marco W. Migliaro
VP, Technical Activities: Fernando Aldana

Developments in multimedia

Video on Silicon Graphics

Tektronix has announced the VideoDesktop software environment and toolset for animation developers using Silicon Graphics workstations.

ChromaMetrix software complements standard video-signal test equipment by analyzing computer-based RGB images for color saturation, frequency, and amplitude prior to recording an image. A time-code calculator simplifies recording by performing mathematical operations on time codes and frame numbers. Both tools come in NTSC and PAL formats.

Image and sequence features let animators view images and define image sequences for subsequent processing or recording. Images are resized to video resolution and interlaced for the NTSC and PAL formats.

VideoDesktop can stand alone or serve as an interface to the Avanzar Video System for Silicon Iris workstations. Avanzar lets workstation users generate studio-quality video output in digital component and digital composite forms in parallel and serial formats. The system also generates RGB, Tetacam, MII, S-Video, and composite NTSC and PAL formats with real-time filtering and encoding.

VideoDesktop costs $1,495, and the Avanzar Video System costs $7,995.

Reader Service Number 35

Multimedia touch monitor

MicroTouch Systems has introduced the TruePoint Touch Monitor for multimedia applications.

The package includes a 14-inch VGA display, an integrated high-resolution capacitive touch screen, controller, and mouse emulation drivers for MS-DOS, Windows, and OS/2. These drivers support the Windows-based MPC multimedia standard and the OS/2-based Ultimedia standard to let users work with Authorware, Icon Author, and Toolbox development tools.

The monitor has a dot pitch of 0.39 mm, a maximum resolution of 640 × 480, and a 30-MHz video bandwidth. It also comes with a tilt/swivel base and video cable.

The $1,399 package works on IBM PCs, PS/2s, and compatibles.

Reader Service Number 36

Videophone allows 3D fax

Colby Systems has announced a videophone that transmits and receives black-and-white still-frame pictures over telephone lines. The ColbyVision videophone consists of a 9-inch high-resolution monitor, an integrated video camera and support arm, and a coder-decoder for digital compression and transmission over the phone. The codec uses noise-immune PCM encoding and transmits with 50 levels of gray.

The system is compatible with standard video systems and PCs, and transmits/receives a 242 × 200-pixel gray-scale still in 12 seconds, according to the company. A CCD camera on an adjustable boom operates at 0.03 foot candles to image objects as close as 2 inches away.

Users can send a "3D fax" and transmit images from video cameras and VCRs, or view them on TV monitors. Images can be saved to computers that use a digitizing card, which allows the system to act as a 3D scanner for paint programs.

The videophone system is priced at $1,499.

Reader Service Number 37

IBM authoring tool

IBM has released the LinkWay Live DOS-based authoring and presentation tool for full-motion video, sound, animation, and high-resolution graphics.

The mouse-driven program lets analog motion video from a videodisc be played through a video window with an overlay card like M-Motion. Digital video captured from other media and stored on the hard drive can run at low speeds without a special card or at 30 frames per second with IBM's ActionMedia II Digital Video Interactive (DVI) card.

Users can select from IBM Speech Adapter, M-ACPA, Digispeech, Sound Blaster, and PS/1 digital audio devices to eliminate rerecording sounds prior to use in presentations. Analog sound support is also provided for CD-ROM applications.

LinkWay Live requires a system with an Intel 286 or higher (386 for DVI) and 2 Mbytes of extended or expanded memory. Networked systems require special software.

The single-system version is $196 for educators and retails at $280. Network licenses range from $980.

Reader Service Number 38

On the wall

RGB Spectrum's MediaWall presentation system integrates computer animation, graphics, and text with full-motion and still video on a wall of monitors or projectors.

In Projector mode, graphics are split across the full array on the wall. Users create effects on their own software; the wall does not require programming. In Monitor-by-Monitor mode, each monitor displays an image from the Macintosh and an auxiliary video source, such as a video recorder, or a combination of both. High-resolution mode lets users tile graphics or still videos of up to 3,200 × 2,400 pixels on individual monitors to form an image exceeding HDTV and 35-mm-film resolutions.

Images are produced through independent horizontal and vertical linear interpolation. Graphics can be split across the display array and zoomed to cover portions or all of the wall.

Reader Service Number 39

Emerging technology: Pen in hand

Graphical environment

The Pen/GEOS graphical operating environment from GeoWorks consists of a pen-oriented user interface, plus GEOS (Graphical Environment Object System), the company's graphical windowing environment for Intel machines. The user interface integrates core pen capabilities: application interaction, handwriting recognition, gestures, and ink capture.

GeoWorks reports that PenGEOS has been demonstrated using handwriting recognition software from Palm Computing. PenGEOS is scheduled for availability to OEMs for shipping late this year.

Reader Service Number 40

Stroke of the pen

Shannon Associates' Coup-de-plume product line, consisting of device drivers and a stand-alone program for evaluating them, integrates free-format, hand-written communication into pen-based applications. The Coup-de-plume device driver for hand-printed characters runs under DOS and uses GTCO and Summagraphics tablet interfaces.

The Evaluator free-standing program emulates driver functions and provides a user interface for evaluating script recognition, adjusting key variables, visualizing the recognition process and derived results, and validating vendor graphic digitizers and pen styluses.

Coup-de-plume/Evaluator sells for $189.

Reader Service Number 41

Stylos speeds the mark

The Stylos/Markup software application for annotating technical graphics documents lets pen computer users open digital graphic documents, zoom in and out of them, and mark directly on the displayed image.

According to Stylos Development Corp., the package will be able to read and write to most industry-standard graphics formats, including AutoCAD's DWG, Intergraph's DGN, TIFF, PCX, BMP, and PICT. It will support Go Corp.'s PenPoint and Microsoft Windows for Pen Computer.

Stylos/Markup is scheduled for release in June at $895, with quantity discounts.

Reader Service Number 42

Desktop to pen-based

CalComp Digitizer Products Group has developed DisplayPad, which lets users turn a standard desktop or portable computer into a pen-based system. DisplayPad uses company electromagnetic digitizer technology in a cordless stylus that features tip-pressure sensing, tilt-angle sensing, tip-height sensing, and graphic-effects controls. The tablet is a high-resolution LCD grid printed on ultrathin transparent film that acts as an antenna for pen-transmitted signals. The tablet features 640 × 480-dot resolution with 64 shades of gray.

The software is compatible with DOS, Macintosh, and Sun systems, and Microsoft Windows for Pen Computing, Go Corp. PenPoint, and CIC PenDOS. It is installed by replacing the existing VGA display card or chip in desktop PCs.

DisplayPad is $2,195.

Reader Service Number 43

Pen communications

Traveling Software has announced Pen LapLink and PenConnect, which incorporate the company's Universal Communications Object module of programming code. UCO can be plugged into software and hardware to provide platform-independent data linking and sharing across hardware systems, operating systems, and data-transport media.

Pen LapLink allows data linking via serial or parallel ports between pen systems and desktop platforms. It uses a specially designed Penbar to display functions and can install itself onto a connected PC. Available to OEMs, Pen LapLink is bundled with NCR's 3125 pen computer, working in concert with NCR's PenOS. It is also compatible with systems using CIC's PenDOS.

PenConnect (also an OEM product) permits disk sharing between MS-DOS desktops and pen computers running PenDOS or Microsoft's Windows for Pen Computing via the serial and parallel ports. PenConnect uses Traveling Software's redirector technology, which makes the drives on one computer available to a second computer.

Reader Service Number 44

Pen applications

Slate Corp. has announced second-quarter availability of three software packages for pen computer systems.

According to the company, the SafetyPen suite of utilities for Go Corp.'s PenPoint operating system provides a multilevel protection scheme that lets users control data access. The suite includes a backup utility for full, partial, and incremental backups, stored in the industry-standard PKZip data-compression format.

Slate is teaming with Day-Timers Inc. to develop Day-Timer PenScheduler software, which includes a multiple-view calendar; a to-do list, an address book, and a notepad.

PenScheduler accepts electronic "ink" exactly as it is drawn on screen. Users can then move, copy, delete, sort, and magnify the information. They can also select part of a page, attach an index category to it, and file it.

PenScheduler works with Microsoft Windows for Pen Computing and the Go PenPoint operating systems, and exports and imports data from desktop applications.

The At-Hand spreadsheet program runs under the PenPoint operating system and gives those systems a 256 × 8,192-cell spreadsheet with over 90 functions. The At-Hand interface is "pencentric": Users can write directly into spreadsheet cells and create formats and reports.

SafetyPen Utilities is $99, Day-Timer PenScheduler $195, and At-Hand $295.

Reader Service Number 45

29th ACM/IEEE DESIGN AUTOMATION CONFERENCE® AND EXHIBITION

DESIGN AUTOMATION 29th DAC CONFERENCE®

ANAHEIM CONVENTION CENTER
June 8 - 12, 1992
Anaheim, California

acm SIGDA® IEEE COMPUTER SOCIETY DATC® IEEE CIRCUITS & SYSTEMS SOCIETY®

NEW DAC USER TRACK

ATTEND THE WORLD'S PREMIER CAD/CAM CONFERENCE!

DA Users and Developers, world-wide attend DAC to keep current in the latest Design Automation Tools and Algorithms.

- Over 130 exhibitors of CAD hardware and software products.
- Over 60 exhibitor technical presentations Monday, June 8, 1992, many announcing new products.
- Monday Exhibits Only Passes. Free invitations available from participating companies.
- Twelve full-day tutorials, Friday, June 12, 1992.
- Over 140 papers, tutorials and panel presentations.

ADVANCE REGISTRATION ENDS MAY 8, 1992!

HOTEL RESERVATION FORM
29th Design Automation Conference
June 8 - 12, 1992

Hotel reservations will be handled by the Anaheim Convention and Visitors Bureau, on a first-come, first-served basis. The following hotels are available:

HOTEL (prioritize your choices)	SINGLE (please circle one)	DOUBLE
___ Convention Center Inn	$ 55.00	$ 55.00
	(Queen Parlor Suites - $78.00)	
___ Anaheim International Inn	62.00	62.00
___ Castle Inn	64.00	64.00
___ Raffles	64.00	64.00
	(Parlor Suites - $79.00)	
___ Holiday Inn	78.00	78.00
___ Quality Inn	79.00	84.00
___ Jolly Roger Inn	79.00	85.00
___ Grand Hotel	82.00	88.00
___ Inn at the Park	85.00	95.00
___ Hyatt Regency Alicante	89.00	95.00
___ Pan Pacific	92.00	98.00
___ Disneyland Hotel	100.00	100.00
___ Anaheim Marriott	110.00	120.00
___ Anaheim Hilton	115.00	130.00

Above rates honored only if your reservation is made by May 8, 1992.

Complete the form below and send to:

DAC HOUSING BUREAU
P.O. Box 4270
Anaheim, CA 92803
(714) 999-8939 - FOR INFORMATION ONLY!

CAUTION:
Hotel reservations sent to the DAC office will be discarded!

Name____
(first) (last)
Company_____________________________Mail Stop______________
Mailing Address____
City ___________________________ State_________ Zip____________
Country_____________________Telephone (____)___________________
Arrival Date__________________ Time__________ AM______ PM______
Departure Date________________ Time__________ AM______ PM______

Reservations will be confirmed by the DAC Housing Bureau and the hotel to which you are assigned. A one night deposit of $80.00 per room is required by check or credit card. Make check payable to the DAC HOUSING BUREAU.

☐ VISA ☐ MASTERCARD ☐ AMERICAN EXPRESS

Number_____________________________Expiration Date___________
Signature____

ADVANCE REGISTRATION FORM
29th Design Automation Conference
June 8 - 12, 1992

To register, mail your payment with this form or a copy to:

29th Design Automation Conference
c/o DAC REGISTRATION DESK
7490 Clubhouse Rd. #102
Boulder, CO 80301 USA

For Information Only, Call (303) 530-4333. **NO TELEPHONE REGISTRATIONS**

IMPORTANT - Registration Instructions

1. Full payment in U.S. Dollars (drawn from a U.S. bank) MUST accompany registration. Any registration without payment WILL be discarded.
2. Advance registrations must be postmarked no later than **May 8, 1992.**
3. Register one person per form (copy this form as needed).
4. **Refund Policy**: Written requests for refunds must be received in the DAC office no later than **May 8, 1992**. Refunds are subject to a $15.00 processing fee.
5. Make checks payable to: **29th Design Automation Conference.**

PLEASE TYPE OR PRINT CLEARLY

CAUTION:
This information will be used to print your ID badge!

Name____
(first) (last)
Company____
Street Address_______________________________________
Mail Stop____
City______________________State________Zip___________
Country_________________Telephone (____)_____________
MEMBER NO. ACM/IEEE__________________________________

CONFERENCE REGISTRATION FEES

	Conference Only	Conference and Tutorial
Advance Registration (postmarked by May 8, 1992)		
Member IEEE/ACM	$140.00	$265.00
Non-Member	$180.00	$355.00
At Conference Registration (after May 8, 1992)		
Member IEEE/ACM	$180.00	$305.00
Non-Member	$220.00	$395.00

Students and One-day-only registrations will be accepted AT CONFERENCE only!

☐ VISA ☐ MASTERCARD

Number_____________________________Expiration Date___________
Signature____

ICM

IC Announcements

Company, Model, Function	Comments	R.S. No.
Anadigics ALD30011 Laser driver	GaAs chip modulates a laser diode at bit rates up to 2.5 Gbits/s. Frequency response ranges from 10 kHz to 3 GHz. On-chip 180° phase splitter enables single-ended input drive, voltage-controlled offset current of 3 to 80 mA, and single –5.2V power supply. Cost (in 100s): $97.	121
Analogy A/D converter models	Behavioral models of 44 commonly used A/D converter components are based on specifications from Analog Devices and Burr-Brown. The models are characterized for such performance parameters as conversion speed, signal noise ratio, and harmonic distortion.	122
Cirrus Logic CL-CD1864 Communications controller	Intelligent 8-channel serial I/O device includes an embedded proprietary RISC processor dedicated to controlling communications functions. Each full-duplex asynchronous communication channel can transfer digital data to a host computer or LAN at rates up to 64 Kbps. Packaged in a 100-pin quad flat pack. Cost (in 1,000s): $33.	123
Digital Equipment 21064-AA RISC microprocessor	This 0.75-µm CMOS-based microprocessor is the first chip to implement Digital's Alpha 64-bit RISC architecture. It offers dual instruction issue and a 150-MHz cycle time. Peak instruction execution is 300 MIPS; supply voltage, 3.3V. Evaluation quantities are available. Cost: $3,375 (in units of 1 to 100); $1,650 (101 to 1,000); and $1,559 (over 1,000).	124
Hitachi America HN28F4001 Flash memory	Fabricated with 0.8-µm CMOS technology, this 4-Mbit memory chip is configured 512K × 8 and comes in versions with maximum access times of 120, 150, and 200 ns. Includes both manual and automatic commands for programming, chip erase, and selective block erase. Typical active power dissipation is 30 mA; standby power dissipation is a maximum of 20 µA. Comes in 32-pin plastic DIPs, TSOPs-Type 1, and SOPs. Cost (in 1,000s): $49.	125
Motorola MC68HC16Y1 Microcontroller	Control unit incorporates CPU based on 16-bit implementation of 68HC11 family. It combines 48 Kbytes of ROM and a RISC-based time-processing unit on chip. TPU has 16 timer channels to handle complex time-critical tasks autonomously. Other features include a general-purpose timer with 11 channels, a 10-bit 8-channel A/D converter, and a multichannel communication interface. Beta-sampling quantities available in a 160-pin QFP. Cost: $38.69.	126
National Semiconductor BSI-2 Interface circuit	Bus system interface circuit for FDDI networks runs at up to 33 MHz with zero wait states and can deliver more than a 760-Mbit burst transfer rate. Incorporates 2 Kbytes of FIFO buffering and allows more than 80-ns bus latency. Other features include direct-connect circuitry for the SBus and a programmable frame preprocessor that sorts frames into queues. The BSI-2 is pin and software compatible with the original BSI. Cost (in 1,000s): $69.	127
Newer Technology Video SIMMs	Single in-line memory modules (256K) for Macintosh Quadra 700 and 900 use 100-ns, dual-ported, surface-mount video RAMs. They operate at 25 MHz on "stock" Quadra and at 33 MHz with New Technology's Quadra/Overdrive accelerator. Cost (set of two): $120.	128
Power Integrations PWR-SMP260 Power supply/battery charger	Universal off-line switching power supply incorporates current-mode controller that contains all blocks required to drive and control power stage. Feed-forward circuit in summing junction maintains constant-power battery charging. DC power outputs up to 60W from rectified 220/240 VAC inputs or 30W from 85-265 VAC inputs. Packaged in 23-pin plastic SIP. Cost (in 1,000s): $4.25.	129
Texas Instruments TLC32047 AIC	Wide-band analog interface circuit is designed to meet 11.4-kHz industry-standard bandwidth for multimedia audio processing. Integrates user-programmable, analog-input, bandpass/low-pass filter with nominal bandwidth of 450 to 11,400 Hz. Maximum sampling rate is 25 KHz. Operates in dual-word, word, and byte modes. Cost (in 1,000s): from $15.33.	130
Trident TLCD9000 Graphics controllers	Super VGA controller for ISA bus systems offers 70-Hz vertical refresh at 1,024 × 768 resolution in 16 colors, noninterlaced. Incorporates feature connector interface logic, data bus transceiver, built-in monitor detection, and automatic memory detection. Supports 256K × 4 DRAMs and interfaces with 256K or 512K memory. Packaged in 160-pin PFP.	131
White Technology WS-128/256/512-K8 CX CMOS SRAMs	Family of 1-Mb (128K × 8), 2-Mb (256K × 8), and 4-Mb (512K × 8) SRAMs in 32-pin DIP with access times ranging from 25 to 120 ns offer 5V operation and CMOS- and TTL-compatible I/O. At 25°C and 5 MHz, typical power consumption ranges from 165 to 825 mW; standby current ranges from 1 to 60 mA; and data retention current, from 1 to 50 µA. Devices are ruggedized for hostile environments. Cost (in 100s): from $90 (128); $135 (256); $203 (512).	132

Microsystem Announcements

Company, Model, Function	Comments	R.S. No.
ATTO Technology SiliconDisk Plus Solid-state disk drive	High-speed 40-MHz SCSI-2-based disk drive for Macintosh, Sun, and PC environments eliminates such delays as seek times and rotational latency. Delivers 0.02-ms access times and 10-Mbyte/s data-transfer rates. Offloads RAM disk functions and uses 32 SIMM sockets to provide up to 512 Mbytes of memory. Cost: $2,995.	141
Cisco Systems CSC-MC+ Flash memory card	Supports flash EPROM circuits in software upgrades. Contains 4 Mbytes of flash EPROM and 64 Kbytes of configuration memory. Typical software load time is 15 to 20 sec.; erase/reprogram times are 2 to 3 min. Available with Cisco's AGS+, MGS, and CGS router/bridges; ASM and MSM terminal servers; and CPT protocol translator. Cost: $2,400.	142
Computer Dynamics SBC-SX$_{LP}$ Single-board computer	Places all functions of a high-speed 386SX-based PC/AT on a single low-power card. Designed for OEM and portable applications that need 100% PC compatibility. Measures 5.75×7.75 in., draws 4.3W of power, and offers 32-bit mathematics. Drives CRTs and a variety of flat panel displays. Cost: from $971 (100s), including 16-MHz CPU and VGA graphics.	143
Cyclone Microsystems CVME962 Single-board computer	VME-based hardware platform built around the 25-MHz i960MC RISC engine. Features 128 Kbits of SRAM, 8 Mbytes of DRAM, up to 1 Mbyte of flash ROM or 4 Mbytes of EPROM, and an 82596 Ethernet coprocessor. Other features include a master/slave VME interface, and software-programmable status LEDs. Ada support is planned. Cost: $8,600.	144
Data Translation DT2839/2896 Channel-count and expander boards	PC/AT-compatible add-on high-channel board and companion channel expander. The DT2839 can acquire data at 1 MHz on a single channel or can be configured to support flexible multichannel scanning at aggregate rates of up to 416 kHz for 32 single-ended or 16 differential input channels. Each DT2896 increases capacity by 96SE/48DI inputs for analog-intensive monitoring applications. Cost: $3,495 (DT2839); $995 (DT2896).	145
Digimetrie PC-DSP 56K-ST Stereo signal-processing board	Fast coprocessor board for IBM XT/ATs and compatibles in audio signal processing and high-fidelity stereo applications. Built around Motorola's DSP 56001, operates at 10 MIPS, with a clock rate of 20 MHz (optional 33 MHz). Standard version comes with 576 Kbytes of RAM. Comes with MIDI and extensive signal processing library.	146
Dover Electronics Manufacturing West ESP 386SX Miniature form factor	Engineering development and evaluation kit measures 5.3×4.1×2 in. Consists of fully integrated PC/AT on four 1.7×5.2 boards: CPU, memory, AT I/O, and power supply. Provides a 286SX processor running at 16, 20, or 25 MHz. Features include on-board Super VGA driver supporting either LCD displays or monitors, power supply, and up to 20 Mbytes of DRAM.	147
Force Computers FBackPlane-4/5/7/14 Backplanes	Four backplanes (4-, 5-, 7-, and 14-slot configurations) meet Futurebus+ specification. Power is bused across the backplane to distribute current load efficiently. Used with FSubRacks. Available in May. Cost: $920 (4-slot); $1,035 (5-slot); $1,355 (7-slot); $2,130 (14-slot).	148
Force Computers FSubRack-5/14 Backplane subracks	FSubRack-5 supports up to 5- slot backplanes; measures 298 mm high × 160 mm wide × 345 mm deep. FSubRack-14 accommodates up to 14-slot backplanes; designed for metric mounts and also fits into existing 19-in. racks. Both versions provide full EMI/RFI shielding on all sides except the front, where Futurebus+ modules are inserted. Available in May. Cost: $195 (FSubRack-5); $495 (FSubRack-14).	149
GrassRoots 2020i Floppy drive subsystem	This 20-Mbyte floppy drive kit includes a 3.5-inch drive suitable for installation in a 3.5-inch drive bay, along with a 5.25-inch adapter frame, a SCSI host adapter, a proprietary formatting program on 720K diskette, a ribbon cable, and one 20-Mbyte diskette. The latter stores up to 40,000 files and is compatible with MS-DOS 3.30, 4.01, 5.0, and Windows 3.0. Cost: $595.	150
Neural Computer Sciences Accelerator card	Plug-in accelerator card for PC/AT compatibles improves performance of NCS's Windows 3.0 neural network package, Neural Desk, by 100 times, according to company sources. Based on a 32-bit DSP, the 25-Mflop card implements the algorithms required for Neural Desk to make over 300,000 neural connections/s. Cost: £1,500.	151
Themis Computers SPARCard 2LC/2SE VMEbus SBCs	Sparcstation 2-compatible boards are rated at 30 MIPS using 40-MHz Sparc processors for integer and floating-point units. Model 2LC targets OEM-embedded applications; includes 64-Kbyte cache, 8-32-Mbyte DRAM, Ethernet, SCSI, two serial channels, and a full IEEE 1014-87 VMEbus interface. Model 2SE additionally includes two standard SBus peripheral controller slots. Cost: $5,995 (2LC); $7,995 (2SE).	152

Supercomputing: The third modality

Ware Myers, Contributing Editor

During Sidney Fernbach's leadership of the Computation Directorate at the Lawrence Livermore National Laboratory, high-performance computing experienced two modalities. It is now on the verge of a third modality, Bill Lokke, LLNL associate director for computations, told an invited audience of 200 at the Sid Fernbach Memorial Symposium: Expanding Computational Horizons. The symposium took place February 28 near Livermore in San Ramon, California, nearly a year after Fernbach passed away (*Computer*, April 1991, p. 74). He died on February 15, 1991.

The Fernbach symposium was the first in a series of events celebrating the 40th anniversary of the lab's founding. Fernbach, employed in September 1952, was one of LLNL's first staff people. For many years, he was also a leading figure in staging the IEEE Computer Society's Compcon Spring event. After retirement, he influenced the direction of high-performance computing as chair of the IEEE Subcommittee on Scientific Supercomputing of the Committee on Communications and Information Policy of the United States Activities Board.

The three modalities. The first modality, back in the 1950s, was simple direct access to the computer, Lokke outlined. This mode was succeeded in the early 1960s by time-sharing. To this day, time-sharing remains the only way to access high-performance computers, though increasingly powerful workstations and networks greatly facilitate this access.

The third modality would be a paradigm change comparable in scope to the change from the first to the second modality. Trying to imagine what it might be is one way of approaching what high-performance computing may become in this decade.

"The third modality would need a variety of characteristics," Lokke said. These characteristics will likely include greatly enhanced versatility, the capability to handle many more users, and distributed software to make better use of locally owned resources.

Increasingly powerful processing elements have opened the door to massively parallel processing, Lokke continued. This surge in capability provides the ability to solve larger problems. After more than a decade of development, massively parallel processing is on the threshold of entering the general-purpose computer market. He expects that within the next year and a half, several vendors will have thousand-processor systems with a hundred megaflops per processor — speaking in terms of orders of magnitude, not precise numbers. With further chip development, capability may increase by another order of magnitude.

LLNL is supporting the development of a high-performance switch to service a gigabit system with a bandwidth of as much as two terabits per second, Lokke reported. The system is being tested.

"Today, the central computing units are saturated, as they always have been," he said. "In the past, we had to use every cycle of a supercomputer, but now a new way of thinking about that is growing. If you are using every cycle, you are employing the machine to full capacity, but you are denying ready access to users."

Farther out are quantitative changes such as optical computing and neural networks. "Neural networks are with us now," he noted, "but it will be a long time before they become commonplace. They could help us to the next great limit — artificial assistance in human tasks. The one I picture is software generation, because all the progress I am talking about hinges on the ability to write programs."

Managing technological change. "We can manage the future only to the extent that we can anticipate it," Jack Worlton quoted Alvin Toffler as saying. Worlton was attached to the Los Alamos National Laboratory from 1954 to 1985 and now heads Worlton Associates of Salt Lake City. "We can anticipate the future only to the extent that we can understand the patterns of change, the principles of change, and the pitfalls that we might fall into along the way," Worlton continued. To this end, Worlton has charted the past progress of the technologies and system-level capabilities that underlie high-performance computing. He flashed several score of these patterns before the audience.

From his studies, he has derived a basic principle: Technologies advance in cycles of qualitative and quantitative change. The shift from electromechanical to electronic computing, for instance, was a qualitative change. That change led, in turn, to a quantitative change of about three orders of magnitude. "Since then we have increased the execution rate by another factor of a million," Worlton stated.

Citing Richard Hamming in a 1965 paper, Worlton argued that "when you change the fundamental parameters of a technology by orders of magnitude, you not only enable the ability to perform the same functions faster and better, but you also often enable the ability to perform entirely new functions. A qualitative change results from the quantitative change. Change of this magnitude has had profound effects, including the way we do science." Worlton's explorations of change patterns revealed six typical relationships:

(1) Linear change over a period of time;

(2) A change in the linear rate from time to time, or piecewise linear, as he called it;

(3) Little change as the final linear piece becomes asymptotic to an ultimate limit, such as one set by the laws of physics; and

(4-6) Exponential change in the same sequence.

For example, the linear trend in the line widths of the Motorola 68000 was almost a straight line from 1982 to 1986. "When I extended the line, it went to zero in 1993," he said. "That is not going to happen!"

In 1986, there was a break in the line-width trend line. The rate of decrease in the line width slowed down. Ultimately, it will be asymptotic to zero, the ultimate limit.

In short, a rapid rate of change, whether linear or exponential, eventually slows down and runs into a limit. The complete curve is S-shaped. During the rapid growth of a new technology, however, we can see only the part of the curve down to date. We may be unrealistically optimistic about the future rate of growth.

Worlton notes a number of pitfalls in trying to predict the course of technology:

- Beware of mistaking short-term change for long-term change. "If you are looking at long-term exponential trends, you ought to have a warning in your mind that they are going to turn into piecewise exponentials," Worlton advised. "When the annual rate of change slows down and turns negative, you know you are at the point of flexure."
- Beware of technological wonder. "We make our greatest mistakes when we get so enamored of a new technology that we forget its disadvantages."
- Beware of rushing a new technology into production prematurely. "Here we have a mature technology that we are using productively. Here we have a number of exploratory technologies vying for our attention, money, and effort. Remember, about 50 of these new ideas fail for every one that succeeds. So be skeptical."
- Beware of the moving target, or "the sailing-ship effect." When the feasibility of steam-powered vessels arose several hundred years ago, sailing-ship builders improved their product to increase their longevity in the marketplace. As such, they delayed the introduction of steam for about 20 years. Mature manufacturers still delay introducing new technologies by potential competitors in this way.

A good exercise in technological assessment is determining where we are in massively parallel computing. "Every conference seems to have a session that asks that question," Worlton observed. In trying to build a framework to systematically study the question, he identified two basic driving forces: the push of the technology and the pull of the market. On the hardware side, the progress in VLSI, which led to RISC processors and inexpensive dynamic RAMs, is among the driving forces. Unfortunately, the software side continues to present problems.

In the long run, however, "the brain is existence proof that we can do massive parallelism on a scale we haven't even thought about," he pointed out.

A high-risk business. "Supercomputing is the most high-risk part of what is already a high-risk industry," asserted Robert Price, former chief executive officer of Control Data, who, beginning in 1953, spent the early part of his career at LLNL. "The reason is, very straightforwardly, that supercomputers involve both extraordinary technical risk and extraordinary market risks."

Price knows whereof he speaks, because Control Data not only built a series of successful supercomputers in the 1960s and 1970s, but also had supercomputer projects that failed.

The technical risk may seem obvious, he went on. Beyond the obvious, however, are the extremely subtle balances that must be struck between architecture and circuit technology. "In oversimplified terms, because the design of a supercomputer takes roughly four years, you are aiming at a point four years out. Trying to anticipate the state of technology that far out can lead and has led to absolute dead ends. On the other hand, working only with proven technology will almost guarantee that you will not successfully implement the most advanced design concepts."

If that were not enough, advanced architectural concepts require new software and new approaches to, and algorithms for, the solution of complex problems, he added. Over all hangs the specter of market risk.

"The most advanced computers are, by definition, addressing a limited market, at least initially," he explained. Those few initial customers are the key to the ultimate success of market development. Moreover, market access is a political as well as an economic matter; this may lead to procurement delays, placing financial stress on the hard-pressed vendor.

"The past is filled with examples of technical and market risk, and failure," he said, reviewing Control Data's record.

In spite of these risks, however, "the development of the supercomputer is one sparkling example of where government and industry worked together to achieve technology pull," Price said. "That wasn't consciously done as a government economic policy, certainly not as industrial policy. It wasn't done as an economic competitiveness initiative, but it was real, it had the appropriate result, and we should learn from it."

Moreover, Fernbach, the man in whose memory the symposium was held, was a major contributor to making it happen. "There have been very few customers like Sid Fernbach who could and would, with great flair, point out the error of your ways, but also the importance of the task," Price concluded.

Design Automation Conference to address latest research topics

Five full days of design automation activities are planned for this year's Design Automation Conference, to be held at the Anaheim Convention Center, Anaheim, California, June 8-12.

The primary goal of the conference is to provide a program that meets the needs of design automation researchers and developers and of engineers who use DA systems to design, test, and manufacture circuits and systems. Topics to be covered range from the most recent technical innovations in research and development to the latest in commercially available systems.

New product introductions and application notes will be highlighted on June 8 in the exhibitor presentations, and more than 115 companies will exhibit the latest commercially available products June 8 through 11.

The technical program consists of papers, panel sessions, and tutorials. The best 125 papers were selected from more than 440 submissions. A major addition this year is the added emphasis on the use of DA tools and systems. The three-day technical program will cover such topics as DA systems, frameworks, design methodology, and manufacturing interfaces.

Four parallel tracks June 9 through 11 will present the latest technical innovations in DA research and development and in the use of DA tools. Nine full-day tutorials will be offered on June 12.

Other topics of interest this year include field-programmable gate arrays; logic, register transfer, and behavioral synthesis techniques; design optimization to increase system performance; and concurrent design of hardware and software.

Expanded panel sessions will cover technology-independent design, ASICs (application-specific integrated circuits), data models and libraries, and CAD tools.

Conference proceedings will be available from the IEEE Computer Society Press in June and can be ordered by phoning (800) CS-BOOKS or, in California, (714) 821-8380. The Order No. is 2822.

For more conference information, contact the Design Automation Conference, 7490 Clubhouse Rd., Ste. 102, Boulder, CO 80301, phone (303) 530-4333, fax (303) 530-4334, e-mail mpa@dac.com.

CALL FOR PAPERS

IEEE International Conference on Wafer-Scale Integration
January 20 - 22, 1993

Fairmont Hotel
San Francisco

Sponsored by the **IEEE Computer Society** and the
IEEE Components, Hybrids, and Manufacturing Technology Society

General Chair : Dr. Peter W. Wyatt MIT Lincoln Laboratory

Program Chair : Prof. R. Mike Lea Brunel University
Uxbridge, England UB8 3PH

The IEEE International Conference on Wafer-Scale Integration is the premier forum for reporting new WSI devices, their architectures and applications, and progress in WSI design, fabrication and packaging technologies. Papers relating to both **Monolithic-WSI** and **Hybrid-WSI** (Multi-Chip Modules) are sought, with equal priority, and preference will be given to reportage of practical contributions to this rapidly expanding field. Exhibition of WSI devices is also strongly encouraged.

Submissions are invited on the following **monolithic-WSI** and **hybrid-WSI** topics:

- applications
- WSI packaging
- WSI yield
- reconfiguration
- CAD for WSI
- I/O techniques
- architectures
- thermal analysis
- test & testability
- rework & self-repair
- cost modelling
- 3-D interconnect
- device & system implementations
- power, clock & signal distribution
- defect & fault tolerance
- WSI device reliability
- cost-effectiveness comparisons
- WSI technology

Authors should submit six copies of paper summaries (3 to 4 pages, plus figures) to the Program Chair. In addition to regular and keynote papers, two evening panel discussions on key WSI issues will also be organised. Regular and keynote papers will be published in a book that will be available at the conference. Suggestions for keynote papers and panel discussions should be sent to the Program Chair before the submission deadline.

Important dates :

June 1, 1992	Summary submission deadline
August 1, 1992	Paper acceptance notification
October 1, 1992	Final manuscript deadline

For further information, contact :

American Chair
Dr. Peter W. Wyatt
Tel. 617 981 7232
Fax. 617 862 9057

European Chair
Prof. R. Mike Lea
Tel. (0)895 203221 (U.K.)
Fax. (0)895 258728 (U.K.)

Asian Chair
Prof. Susumu Horiguchi
Tel. 22 222 1800 (Japan)
Fax.22 263 9418 (Japan)

CALL FOR PAPERS

ACM SIGSoft 92, Fifth Symp. on Software Development Environments: Dec. 9-11, 1992, Washington, DC. Submit six copies of abstract and full paper (maximum length 6,000 words) by **Apr. 21, 1992**, and camera-ready paper by **Sept. 1, 1992**, to Herbert Weber, Fachbereich Informatik, Univ. of Dortmund, Baroperstrasse 301, 4600 Dortmund 50, Germany.

KR 92, Third Int'l Conf. on Principles of Knowledge Representation and Reasoning: Oct. 26-29, 1992, Cambridge, Mass. Sponsors: Am. Assoc. for Artificial Intelligence et al. For submittal-requirement details, contact William Swartout, Information Sciences Inst., Univ. of Southern California, 4676 Admiralty Way, Marina del Rey, CA 90292-6695, phone (310) 822-1511, fax (310) 823-6714, e-mail swartout@isi.edu; or Bernhard Nebel, DFKI, Stuhlstazenhausweg 3, D-W-6600 Saarbruecken, Germany, phone 49 (681) 302-5254, fax 49 (681) 302-5341, e-mail nebel@dfki.uni-sb.de. Submittal deadline: **Apr. 21, 1992**.

PDC 92, 1992 Participatory Design Conf.: Nov. 6-7, 1992, Cambridge, Mass. Sponsor: Computer Professionals for Social Responsibility. Submit paper (maximum 4,000 words) by **Apr. 22, 1992**, and poster abstracts (1,500 words) by **Aug. 15, 1992**, to Michael J. Muller, Bellcore RRC-1H229, 444 Hoes Lane, Piscataway, NJ 08854, phone (908) 699-4892, fax (908) 336-2932, e-mail michael@bellcore.com.

Second Workshop on Management of Replicated Data: Nov. 12-13, 1992, Monterey, Calif. Sponsor: IEEE Computer Soc. Technical Committee on Operating Systems and Application Environments. Submit 12 copies of position statement (2 to 4 double-spaced pages) by **Apr. 24, 1992**, to Hector Garcia-Molina, Computer Science Dept., Stanford Univ., Stanford, CA 94305, phone (415) 723-0685, e-mail hector@cs.stanford.edu.

AICS 92, Fourth Irish Conf. on Artificial Intelligence and Cognitive Science: Sept. 10-11, 1992, Limerick, Ireland. For submittal-requirement details and other information, contact Kevin Ryan, Computer Science and Information Systems Dept., Univ. of Limerick, Ireland, phone 353 (61) 333-644, fax 353 (61) 330316, e-mail aics92@ul.ie. Submittal deadline: **Apr. 24, 1992**.

EWSPT 92, Second European Workshop on Software Process Tech.: Sept. 7-8, 1992, Trondheim, Norway. Cosponsors: Assoc. Francaise pour la Cybernetique Economique et Technique et al. Submit four copies of full paper (6,000-word maximum) by **Apr. 25, 1992**, and camera-ready paper by **June 28, 1992**, to Jean-Claude Derniame, Centre de Recherche en Informatique de Nancy, Campus scientifique B.P. 239, F-54506 Vandoeurve Les Nancy, France, phone 33 (83) 413-052, fax 33 (83) 413-079, e-mail derniame@loria.crin.fr.

Crypto 92: Aug. 16-20, 1992, Santa Barbara, Calif. Sponsor: Int'l Assoc. of Cryptologic Research. Submit 12 copies of detailed abstract (not a full paper) by **Apr. 27, 1992**, and revised version by **July 20, 1992**, to Ernest F. Brickell, Div. 1423, Sandia Nat'l Labs, Albuquerque, NM 87185, phone (505) 845-7655, fax (505) 845-7442, e-mail efbrick@cs.sandia.gov.

ICARCV 92, Second Int'l Conf. on Automation, Robotics, and Computer Vision: Sept. 15-18, 1992, Singapore. Cosponsors: Singapore Institution of Engineers et al. Submit four copies of 300- to 500-word summary on neural networks in engineering and scientific applications by **Apr. 30, 1992**, and final manuscript by **June 30, 1992**, to Yoshiyasu Takefuji, Electrical Eng. Dept., Case Western Reserve Univ., Cleveland, OH 44106, phone (216) 368-6430, fax (216) 368-2668, e-mail takefuji@axon.eeap.cwru.edu; and submit four copies on other subjects to ICARCV Secretariat, Associated Conventions and Exhibitions, 204 Bukit Timah Rd., #04-00, Boon Liew Bldg., Singapore 0922, phone (65) 799-5470, fax (65) 791-2687, e-mail emital@ntuvax.bitnet.

Call for articles and referees for *Computer*

Computer seeks articles for inclusion in a future special issue.

Multichip modules has been selected as the theme for the **April 1993** issue. Manuscripts reporting survey, original research, design and development, and applications of MCM technology are sought immediately. See p. 101 of the December 1991 issue for more details.

Fourteen copies of the full manuscript are due by **June 15, 1992**; notification of decisions is set for **Oct. 15, 1992**; and the deadline for submittal of the final version of each manuscript is **Dec. 1, 1992**.

Submissions and questions should be directed to P.R. Mukund, Dept. of Electrical Eng., Rochester Inst. of Tech., 1 Lomb Memorial Dr., Rochester, NY 14623, phone (716) 475-2174, e-mail prmee@ritvax.isc.rit.edu; or J.F. McDonald, Dept. of Electrical Eng., Rensselaer Polytechnic Inst., Troy, NY 12180, phone (518) 276-2919, e-mail mcdonald@unix.cie.rip.edu.

For submittal to Computer, *manuscripts must not have been previously published or currently submitted for publication elsewhere. Each manuscript should be no more than 32 typewritten, double-spaced, single-sided pages long, including all text, figures, and references. Each of the 14 copies submitted should include a page that contains the title of the article, the full name(s) and affiliation(s) of the author(s), complete postal and electronic address(es) of all the authors as well as their telephone and fax number(s), a 300-word abstract, and a list of keywords identifying the central issues of the manuscript's contents. The final manuscript should be approximately 8,000 words in length and contain no more than 12 references.*

If you are willing to review articles for any special issue, please send a note listing your research interests to Jon T. Butler, editor-in-chief of Computer, *at the Dept. of Electrical and Computer Eng., Naval Postgraduate School, Code EC/Bu, Monterey, CA 93943-5004, Internet butler@ece.nps.navy.mil.*

IEEE Trans. on Reliability plans a special issue in December 1992 on fault-tolerant software. Submit six copies of full manuscript by **Apr. 30, 1992**, to John F. Meyer, EECS Dept., Univ. of Michigan, Ann Arbor, MI 48109, phone (313) 763-0073, fax (313) 763-1503; or Hoang Pham, Idaho Nat'l Eng. Lab, PO Box 1625, MS 2406, Idaho Falls, ID 83415, phone (208) 526-9274, fax (208) 526-2930.

ISAI, Fifth Int'l Symp. on Artificial Intelligence: Dec. 7-11, 1992, Cancun, Mexico. Sponsor: Inst. Tecnologico y de Estudios Superiores de Monterrey. Submit manuscript by **Apr. 30, 1992**, and camera-ready copy by **July 31, 1992**, to Hugo Terashima, ITESM, Centro de Inteligencia Artificial, Sucursal de Correos "J", Monterrey, N.L. 64849, Mexico, phone 52 (83) 58-2000, ext. 5134, fax 52 (83) 58-1400, e-mail isai@tecmtyvm.bitnet or terashim@mtecv2.mty.itesm.mx.

1992 Complex Systems Eng. Synthesis and Assessment Tech. Workshop: July 20-24, 1992, Silver Spring, Md. Cosponsors: Naval Surface Warfare Center Dahlgren Div., Office of Naval Tech. Submit five copies of 7,000-word (maximum) paper by **Apr. 30, 1992**, and camera-ready version by **June 15, 1992**, to Steve Howell, NSWCWODET, Code U33, 10901 New Hampshire Ave., Silver Spring, MD 20903-5000, phone (301) 394-3987, fax (301) 394-1175, e-mail showell@nswc-wo.navy.mil.

ISCIS 7, Seventh Int'l Symp. on Computer and Information Sciences: Nov. 2-4, 1992, Antalya, Turkey. Cosponsors: IEEE Computer Soc. Turkey Chapter et al. Submit three copies of 10-page paper or four-page short communication by **May 1, 1992**, to Erol Gelenbe, Ecole des Hautes Etudes en Informatique, 45 rue des Saints-Peres, 75006 Paris, France, phone 33 (1) 4286-2136, fax 33 (1) 4286-2231, e-mail erol@ehei.ehci.fr.

RTSS 92, 13th IEEE Real-Time Systems Symp.: Dec. 2-4, 1992, Phoenix, Ariz. Sponsor: IEEE Computer Soc. Technical Committee on Real-Time Computing. Transmit synopsis or 80-word abstract by **May 1, 1992,** to rtss92@cis.upenn.edu and submit five copies of paper (5,000-word maximum) to Susan Davidson, Computer and Information Science Dept., Univ. of Pennsylvania, Philadelphia, PA 19104-6389. Submit camera-ready paper to Davidson by **Sept. 15, 1992.**

HPDC 1, First Int'l Symp. on High-Performance Distributed Computing: Sept. 9-10, 1992, Syracuse, N.Y. Cosponsor: Syracuse Univ. Submit five copies of manuscript (maximum 20 double-spaced pages) by **May 1, 1992**, and camera-ready version by **July 10, 1992**, To Salim Hariri, Electrical and Computer Eng. Dept., Syracuse Univ., 111 Link Hall, Syracuse, NY 13244, phone (315) 443-4282, e-mail hariri@cat.syr.edu.

Sixth Int'l Conf. on Modeling Techniques and Tools for Computer Performance Modeling: Sept. 16-18, 1992, Edinburgh, Scotland. Sponsors: Int'l Federation for Information Processing Working Group 7.3, British Computer Soc. Submit complete paper by **May 1, 1992**, to Margaret Davis, Edinburgh Univ., Edinburgh EH9 3JZ, Scotland, phone 44 (31) 650-5128, fax 44 (31) 667-7209, e-mail mda@cs.ed.ac.uk.

Int'l J. of Expert Systems with Applications plans a special issue in 1993 on the blackboard paradigm and its applications. Submit five copies of paper by **May 1, 1992**, and final version by **Aug. 15, 1992**, to Ajay Vinze or Arun Sen, Business Analysis and Research Dept., College of Business Admin., Texas A&M Univ., College Station, TX 77843-4217, phone (409) 845-1616, fax (409) 845-5653, e-mail vinze@tamcba.bitnet or sen@tamcba.bitnet.

ICM 92, Int'l Conf. on Microelectronics: Dec. 19-21, 1992, Monastir, Tunisia. Submit three copies of 35-word abstract and 500-word summary by **May 1, 1992**, and final four-page extended abstract by **Sept. 15, 1992**, to M.I. Elmasry, VLSI Group, Univ. of Waterloo, Waterloo, Ont., Canada N2L 3G1, phone (519) 885-1211, ext. 3753, fax (519) 746-5195, e-mail elmasry@vlsi.uwaterloo.ca (for North Am. and Europe); or R. Tourki, Faculte des Sciences, Route de Kaironan, 5000 Monastir, Tunisia, phone (216) 3-61-766, fax (216) 3-62-873 (for all other areas).

Fifth Digital Signal Processing Workshop: Sept. 13-16, 1992, Starved Rock State Park, Ill. Sponsor: IEEE Signal Processing Soc. Submit four copies of two-page summary by **May 1, 1992**, to Robert A. Gabel, Rm. M-203, MIT Lincoln Lab, Lexington, MA 02173-0073.

DFT 92, 1992 IEEE Int'l Workshop on Defect and Fault Tolerance in VLSI Systems: Nov. 4-6, 1992, Dallas. Submit 20 copies of summary by **May 1, 1992**, and final 10-page manuscript by **Aug. 15, 1992**, to Fabrizio Lombardi, Computer Science Dept., Texas A&M Univ., College Station, TX 77843-3112, phone (409) 845-6841, fax (409) 847-8578, e-mail lombardi@cs.tamu.edu.

IASTED, Int'l Conf. on Reliability, Quality Control, and Risk Assessment: Nov. 4-6, 1992, Washington, DC. Sponsors: Int'l Assoc. of Science and Tech. for Development, IEEE Reliability Soc. Submit three copies of full paper by **May 1, 1992**, and camera-ready version by **Sept. 15, 1992**, to Hoang Pham, Idaho Nat'l Eng. Lab, PO Box 1625, MS 2406, Idaho Falls, ID 83415, phone (208) 526-9274, fax (208) 526-2930.

IEEE Potentials magazine seeks tutorial papers on various aspects of electrical/computer engineering and computer science that would interest students. Submit synopsis by **May 1, 1992**, to George W. Zobrist, Computer Science Dept., Univ. of Missouri-Rolla, Rolla, MO 65401, phone (314) 341-4836, e-mail c2816@umrvmb.umr.edu.

4CCCG, 4th Canadian Conf. on Computational Geometry: Aug. 10-14, 1992, St. John's, Nfld., Canada. Submit seven copies of six-page abstract by **May 1, 1992**, to Cao An Wang, Computer Science Dept., Memorial Univ. of Newfoundland, St. John's, Nfld., Canada A1C 5S7, phone (709) 737-4590, fax (709) 737-2009, e-mail 4cccg@cs.mun.edu.

Technology 2002, Third Nat'l Technology Transfer Conf.: Dec. 1-3, 1992, Baltimore. Cosponsors: NASA et al. Submit 1- to -1.5 page abstract by **May 1, 1992**, to Leonard A. Ault, Code CU, NASA Headquarters, 600 Independence Ave. SW, Washington, DC 20546, phone (703) 557-5598, fax (703) 557-8186.

NAFIPS 92, North Am. Fuzzy Information Processing Soc. Conf.: Dec. 15-17, 1992, Puerto Vallarta, Mexico. Cosponsors: NASA et al. Submit four copies of paper by **May 1, 1992**, and final version by **Oct. 1, 1992**, to Carla Armstrong, Software Tech. Branch/PT4, NASA Lyndon B. Johnson Space Center, Houston, TX 77058, phone (713) 483-9071, fax (713) 438-5200, e-mail carmstrong@gothamcity.jsc.nasa.gov.

UIST 92, Fifth Symp. on User Interface Software and Tech.: Nov. 15-18, 1992, Monterey, Calif. Sponsor: ACM. Submit paper by **May 1, 1992**, to Mark Green, Computer Science Dept., Univ. of Alberta, Edmonton, Alta., Canada T6G 2H1, phone (403) 492-4584, fax (403) 492-1071, e-mail mark@cs.ualberta.ca.

BMVC 92, British Machine Vision Conf.: Sept. 21-24, 1992, Leeds, UK. Sponsor: British Machine Vision Assoc. Submit paper (5,000-word maximum) by **May 1, 1992**, to David Hogg, School of Computer Studies, Univ. of Leeds, Leeds LS2 9JT, UK, phone 44 (532) 335-765, fax 44 (532) 335-468, e-mail dch@dcs.leeds.ac.uk.

Int'l Workshop on Hardware-Software Codesign: Sept. 30-Oct. 2, 1992, Estes Park, Colo. Cosponsors: IEEE Computer Soc., IEEE Circuits and Systems Soc. Submit paper by **May 4, 1992**, to Wayne Wolf, Electrical Eng. Dept., Princeton Univ., Princeton, NJ 08544, phone (609) 258-1424, e-mail wolf@princeton.edu.

IWOOOS 92, Int'l Workshop on Object Orientation in Operating Systems: Sept. 24-25, 1992, Paris. Sponsor: Institut National de Recherche en Informatique et en Automatique. Submit full paper (6,000 words) or position paper (2,000 words) by **May 8, 1992**, to Roy Campbell, Univ. of Illinois, Computer Science Dept., 2413 Digital Lab, 1304 W. Springfield, Ave., Urbana, IL 61801, phone (217) 333-3328, e-mail roy@cs.uiuc.edu or eric.diku.dk.

IEEE Computer Graphics & Applications plans a special issue in November 1992 on semiconductor technology for graphics and imaging. Submit five copies of complete manuscript (approximately 5,000 words) by **May 15, 1992**, to Jack Bresenham, Computer Science Dept., Thurmond Hall, Winthrop College, Rock Hill, SC 29733; Nick England, Visualization Tech., Sun Microsystems, PO Box 13447, Research Triangle Park, NC 27709, or Karl Guttag, Texas Instruments, 12203 SW Freeway, Houston, TX 77251.

ACSAC 92, Eighth Computer Security Applications Conf.: Nov. 30-Dec. 4, 1992, San Antonio, Texas. Cosponsors: Aerospace Computer Security Associates et al. Submit paper by **May 15, 1992**, and camera-ready copy by **Sept. 18, 1992,** to Ann Marmor-Squires, TRW Systems Div., 1 Federal Systems Park Dr., Fairfax, VA 22033, phone (703) 803-5503, e-mail marmor@charm.isi.edu.

SimTec 92, Int'l Simulation Tech. Conf.: Nov. 4-6, 1992, Houston. Sponsors: Soc. for Computer Simulation et al. Submit five copies of full paper by **May 15, 1992**, to Mary Lou Padgett, Auburn Univ., 1165 Owens Rd., Auburn, AL 36830, phone (205) 821-2472, fax (619) 277-3930, e-mail mpadgett@eng.auburn.edu.

First Int'l Conf. on Information and Knowledge Management: Nov. 8-11, 1992, Baltimore. Sponsor: Int'l Soc. of Mini- and Microcomputers. Submit five copies of paper by **May 31, 1992,** and camera-ready paper by **Aug. 10, 1992,** to Yelena Yesha, Computer Science Dept., Univ. of Maryland, 5401 Wilkens Ave., Baltimore, MD 21228-5398, phone (410) 455-3000, fax (410) 455-3969, e-mail cikm@cs.umbc.edu.

Third Eurographics Workshop on Object-Oriented Graphics: Oct. 28-30, 1992, Champery, Switzerland. Sponsor: Eurographics, Univ. of Geneva. Submit four copies of paper by **May 31, 1992**, and revised paper by **Jan. 15, 1993**, to Vicki de Mey or Xavier Pintado, Centre Universitaire d'Informatique, Univ. of Geneva, 12 rue du Lac, 1207 Geneva, Switzerland, phone 41 (22) 787-6586, fax 41 (22) 735-3905, e-mail eoog@cui.unige.ch.

IEEE Software seeks articles for a special issue during the fourth quarter of 1992 on object-oriented software development. Articles addressing the practical aspects of the conceptual, cultural, and technological transition to this new software development paradigm, such as studies and experience reports, are particularly welcome. Submit seven copies of article by **June 1, 1992,** to Wojtek Kozaczynski, Center for Strategic Tech. Research, Andersen Consulting, 100 S. Wacker Dr., Chicago, IL 60606, fax (312) 507-3526, e-mail wojtek@andersen.com; or Annie Kuntzmann-Combelles, Corelis Technologie, 31, av. du General Leclerc, F-92340 Bourg-la-Reine, France, fax 33 (1) 4664-7739. For author guidelines, contact Angela Burgess, phone (714) 821-8380 or e-mail a.burgess@compmail.com.

IEEE Trans. on Knowledge and Data Eng. plans a special issue on multimedia information systems. Submit eight copies of manuscript by **June 1, 1992**, and final revised version by **Nov. 15, 1992**, to Rajiv Mehrotra, Center for Robotics and Manufacturing Systems, Univ. of Kentucky, Lexington, KY 40506, phone (606) 257-6262, ext. 215, e-mail rajiv@s.ms.uky.edu.

IEEE Int'l Symp. on Requirements Eng.: Jan. 4-6, 1993, San Diego, Calif. Submit five copies of paper by **June 1, 1992**, and camera-ready version by **Sept. 15, 1992**, to Anthony Finkelstein, Imperial College, Computing Dept., 180 Queens Gate, London SW7 2BZ, UK, phone 44 (71) 589-5111, e-mail acwf@doc.ic.ac.uk.

VLSI Design 93, Sixth Int'l Conf. on VLSI Design: Jan. 3-6, 1993, Bombay, India. Sponsor: VLSI Soc. of India. Submit eight copies of full paper by **June 1, 1992**, and camera-ready manuscript by **Sept. 15, 1992,** to Srimat Chakradhar, NEC Research Inst., 4 Independence Way, Princeton, NJ 08540, phone (609) 951-2962, fax (609) 951-2499, e-mail chak@research.nj.nec.com or chak@research.nec.com; or Sunil Sherlekar, Computer Science and Eng. Dept., Indian Inst. of Tech., Powai, Bombay 400076, India, phone 91 (22) 578-5708, fax 91 (22) 578-3480, e-mail sds@dhruv.ernet.in or sds@cse.iitb.ernet.in.

WACV, IEEE Workshop on Applications of Computer Vision: Nov. 30-Dec. 2, 1992, Palm Springs, Calif. Submit four copies of paper (30 double-space pages maximum) by **June 1, 1992,** and final paper by **Sept. 1, 1992,** to Martin Herman, Sensory Intelligence Group, Nat'l Inst. of Standards and Tech., Bldg. 220, Rm B124, Gaithersburg, MC 20899.

IEEE Computer Graphics & Applications plans a special issue on applications of advanced graphics in environmental decision-support systems, environmental modeling, natural resource mapping and monitoring, and geographic information systems. Submit article (approximately 5,000 words) by **June 1, 1992,** to Philip K. Robertson and David J. Abel, CSIRO Div. of Information Tech., Center for Spatial Information Systems, GPO Box 664, Canberra, ACT 2601, Australia.

26th Asilomar Conf. on Signals, Systems, and Computers: Oct. 26-28, 1992, Pacific Grove, Calif. Cosponsors: Naval Postgraduate School et al. Submit 50- to 100-word abstract and an extended summary by **June 1, 1992**, to James A. Ritcey, Electrical Eng. Dept., FT-10, Univ. of Washington, Seattle, WA 98195, phone (206) 543-4702, fax (206) 543-3842, e-mail ritcey@ee.washington.edu.

ISIT 93, 1993 IEEE Int'l Symp. on Information Theory: Jan. 17-22, 1993, San Antonio, Texas. Sponsor: IEEE Information Theory Soc. Submit three copies of long paper (complete paper and 180-word abstract) by **June 1, 1992**, and short paper (500-word summary and 180-word abstract) by **July 1, 1992**, to Richard E. Blahut, Mail Drop 0600, IBM, Route 17C, Owego, NY 13827, phone (607) 751-2217.

Advanced Study Inst. on Real-Time Computing: Oct. 5-18, 1992, Sint Maarten/St. Martin, Dutch/French Antilles. Sponsors: NATO et al. Submit manuscript by **June 1, 1992,** to Alexander D. Stoyenko, Computer and Information Science Dept., New Jersey Inst. of Tech., University Heights, Newark, NJ 07102, phone 596-5765, fax (201) 596-5777, e-mail alex@vienna.njit.edu.

ASM 92, Third Int'l Conf. on Applications of Software Measurement: Nov. 16-19, 1992, San Diego, Calif. Sponsors: Data Processing Management Assoc. et al. Submit one- to three-page proposal by **June 1, 1992**, to David Gelperin, Software Quality Eng., 2738 Winnetka Ave. N, New Hope, MN 55427, fax (612) 591-1534.

Compugraphics 92, Second Int'l Conf. on Computational Graphics and Visualization Techniques: Dec. 14-18, 1992, Lisbon, Portugal. Sponsors: ACM et al. Submit five copies of extended abstract by **June 1, 1992**, to Harold P. Santo, Civil Eng. Dept., IST-Advanced Technical Inst., Technical Univ. of Lisbon, Av. Rovisco Pais, 1096 Lisboa Codex, Portugal, phone 351 (1) 847-3421, fax 351 (1) 89-7650.

J. of Parallel and Distributed Computing plans a special issue on parallel and distributed simulation. Publisher: Academic Press. Submit five copies of full paper by **June 1, 1992**, to Jason Yi-Bing Lin, Bellcore, 445 South St., 2D-297, Morristown, NJ 07962-1910, phone (201) 829-5095, fax (201) 829-5884, e-mail liny@thumper.bellcore.com.

DX 92, Third Int'l Workshop on Principles of Diagnosis: Oct. 11-14, 1992, Seattle. Submit three hardcopies of paper (5,000-word maximum length) by **June 2, 1992**, to Ethan Scarl, MS 7L-64, Boeing Computer Services, PO Box 24346, Seattle, WA 98124-0346 (not for courier delivery), phone (206) 865-3255, fax (206) 865-2964, e-mail ethan@atc.boeing.com; or Ethan Scarl, MS 7L-64, Boeing Computer Services 2760, 160th Ave. SE, Bldg. 33-07, Bellevue, WA 98008 (not for US mail delivery).

IEEE Computer Graphics & Applications plans a special issue in March 1993 on amazing desktop applications. Submit five copies of article by **June 15, 1992**, to Jack Grimes, Apple Computer, 20525 Mariani Ave., Cupertino, CA 95014, phone (408) 862-6264, e-mail jgrimes@cup.portal.com; or Mike Potel, Apple Computer, 20525 Mariani Ave., Cupertino, CA 95014, phone (408) 974-6253, e-mail potel@apple.com.

PDIS, Second Int'l Conf. on Parallel and Distributed Information Systems: Jan. 20-23, 1993, San Diego, Calif. Cosponsor: ACM. Submit six copies of paper (8,000 words or 25 double-spaced pages maximum) by **June 15, 1992**, and camera-ready copy by **Oct. 12, 1992**, to

Michael Carey, Computer Sciences Dept., Univ. of Wisconsin, 1210 W. Dayton St., Madison, WI 53706; or Patrick Valduriez, INRIA, Rocquencourt, 78153 Le Chesnay Cedex, France.

Ninth Israeli Conf. on Artificial Intelligence and Computer Vision: Dec. 28-29, 1992, Tel Aviv. Sponsor: Israeli Information Processing Assoc. Submit four copies of full paper by **June 15, 1992**, to Shimon Edelman, Applied Math and Computer Science Dept., Weizmann Inst. of Science, Rehovot 76100, Israel, phone 972 (8) 344-122, fax 972 (8) 342-856, e-mail edelman@wisdom.weizmann.ac.il (for papers on vision); or Ehud Gudes, Computer Science Dept., Ben Gurion Univ., Beer Sheva 84105, Israel, phone 972 (57) 461-626, e-mail ehud@bengus.bitnet (for AI papers).

ISADS 93, 1993 Int'l Symp. on Autonomous Decentralized Systems: Mar. 30-Apr. 1, 1993, Kawasaki, Japan. Cosponsors: Information Processing Soc. of Japan et al. Submit eight copies of full paper by **July 1, 1992**, and final version by **Dec. 29, 1992,** to Domenico Ferrari, Electrical Eng. and Computer Sciences Dept., Univ. of California, Berkeley, CA 94720. For program information, fax to (510) 642-5775, e-mail isads@tenet.berkeley.edu.

IEEE Infocom 93, 12th Conf. on Computer Comm.: Mar. 28-Apr. 1, 1993, San Francisco. Cosponsor: IEEE Comm. Soc. Submit five double-spaced copies of manuscript by **July 15, 1992**, and camera-ready copy by **Dec. 7, 1992**, to Kenneth Vastola, ECSE Dept., Rensselaer Polytechnic Inst., Troy, NY 12180-3590, phone (518) 276-6074, e-mail infocom@ecse.rpi.edu.

RIDE-IMS 93, Third Int'l Workshop on Research Issues in Data Eng. — Interoperability in Multidatabase Systems: Apr. 18-20, 1993, Vienna. Cosponsor: Austrian Computer Soc. Submit six copies of paper (20-page maximum) by **July 15, 1992**, and camera-ready version by **Jan. 1, 1993**, to Amit Sheth, Bellcore, RRC 1J-210, 444 Hoes Lane, Piscataway, NJ 08854-4182, phone (908) 699-3300, fax (908) 336-2969, e-mail amit@ctt.bellcore.com (in the Americas, Asia, and Australia); or Hans-J. Schek, Computer Science Dept., Swiss Federal Inst. of Tech., ETH Zentrum, CH-8092 Zurich, Switzerland, phone 41 (1) 254-7240, fax 41 (1) 262-3973, e-mail schek@inf.ethz.ch (in Europe and Africa).

Moscots 93, Int'l Workshop on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems: Jan. 17-20, 1993, La Jolla, Calif. Sponsors: Soc. for Computer Simulation et al. Submit four copies of paper (10 double-spaced pages) by **July 15, 1992**, to Jean Walrand, Electrical Eng. and Computer Sciences Dept., 267M Cory Hall, Univ. of California, Berkeley, CA 94720, phone (510) 642-1529, fax (510) 643-8426, e-mail wlr@diva.berkeley.edu.

CALENDAR

In the accompanying Calendar and adjoining Call for Papers, the IEEE Computer Society logo identifies the conferences the society is sponsoring or cooperating in. Other conferences of interest to our readers are also listed. The notices are published free of charge, in chronological order, and as space permits.

For inclusion in the Call for Papers section, please submit the name of the event, date(s), location, sponsor(s), deadline for submittals, phone and fax numbers, and — if applicable — the electronic address of the person to whom papers should be directed.

For the Calendar section, please provide the event name, date(s), location, sponsor(s), phone and fax numbers, and the electronic address of the person to contact for complete information.

Submittals should arrive at *Computer* **at least five weeks before the month of publication** (i.e., for the **June 1992** issue, send information for receipt by **April 20, 1992**) to Chuck Governale, Calendar Dept., *Computer*, PO Box 3014, Los Alamitos, CA 90720-1264, fax (714) 821-4010, e-mail c.governale@compmail.com.

April 1992

NATUG, North Am. Transputer Users Group 1992 Spring Meeting, Apr. 19-21, Baltimore. Contact David L. Fielding, Advanced Computing Research Inst., Cornell Univ., 705 Theory Center Bldg., Ithaca, NY 14853-3801, phone (607) 254-8872, fax (607) 254-8888.

ICCL 92, Int'l Conf. on Computer Languages, Apr. 20-23, San Francisco. Sponsor: IEEE Computer Soc. Technical Committee on Computer Languages. Contact Mario Barbacci, Software Eng. Inst., Carnegie Mellon Univ., Pittsburgh, PA 15213, phone (412) 268-7704, fax (412) 268-5758, e-mail barbacci@sei.cmu.edu.

1992 Symp. on Optical Eng. and Photonics in Aerospace Sensing, Apr. 20-24, Orlando, Fla. Sponsor: Int'l Soc. for Optical Eng. Contact SPIE, PO Box 10, Bellingham, WA 98227-0010, phone (206) 676-3290, fax (206) 647-1445; SPIE, Xantener Strasse 22, D-1000 Berlin 15, Germany, phone 49 (30) 881-5047; or SPIE, c/o OTO Research, Takeuchi Bldg., 1-34-12 Takatanobaba, Shinjuku-ku, Tokyo 160, Japan, phone 81 (03) 3208-7821, fax 81 (03) 3200-2889.

15th IEEE Workshop on Design for Testability, Apr. 21-24, Vail, Colo. Contact T.W. Williams, IBM, PO Box 1900, 67A/021, Boulder, CO 80301-9191, phone (303) 924-7692.

Supercomputing Japan 92, Apr. 22-24, Yokohama, Japan. Sponsor: Pacific Meridian Group. Contact MPG, 116 E. Blithedale Ave., Suite 2, Mill Valley, CA 94941, phone (415) 381-2255.

Seventh Symp. on Microcomputer and Microprocessor Applications, Apr. 22-24, Budapest, Hungary. Cosponsors: Scientific Soc. for Telecommunication et al. Contact SST, Budapest 5. Pf. 451., Hungary, H-1372, phone (361) 153-1027, fax (361) 153-0451.

Third Workshop on Workstation Operating Systems, Apr. 23-24, Key Biscayne, Fla. Sponsor: IEEE Computer Soc. Technical Committee on Operating Systems and Applications Environments. Contact Joseph Boykin, 7 Hampton Rd., Natick, MA 01760, phone (508) 651-8228.

SHPCC 92, Scalable High-Performance Computing Conf., Apr. 26-29, Williamsburg, Va. Contact Emily Todd, ICASE, MS 312-C, NASA Langley Research Center, Hampton, VA 23665, e-mail emily@icase.edu; or IEEE Computer Soc., 1730 Massachusetts Ave. NW, Washington, DC 20036-1903, phone (202) 371-1013, fax (202) 728-0884.

C plus C++...In Action, Apr. 27-May 1, Teaneck, N.J. Sponsor: Wang Inst. of Boston Univ. Contact Jacqueline E. Rawlings, Boston Univ., Corporate Education Center, 72 Tyng Rd., Tyngsboro, MA 01879, phone (508) 649-9731, ext. 240, fax (508) 649-6926.

Downsizing/Rightsizing Corporate Computing Conf., Apr. 27-May 1, Los Angeles. Sponsor: Boston Univ. Corporate Education Center. Contact Corporate Education Center, Boston Univ., Tyng Rd., Tyngsboro, MA 01879.

Third Multimedia Conf. on CD-I, Apr. 28-29, London. Contact Event Organisation Co., 8 Cotswold Mews, Battersea Sq., London SW11 3RA, UK, phone 44 (0) 71-228-8034, fax 44 (0) 71-924-1790.

23rd Pittsburgh Conf. on Modeling and Simulation, Apr. 30-May 1, Pittsburgh. Cosponsors: Univ. of Pittsburgh, IEEE et al.

Contact William G. Vogt or Marlin H. Mickle, 348 Benedum Eng. Hall, Univ. of Pittsburgh, Pittsburgh, PA 15261.

15th Nat'l Conf. of the Independent Computer Consultants Assoc., Apr. 30-May 2, St. Louis. Contact ICCA, 933 Gardenview Office Pkwy., St. Louis, MO 63141, phone (314) 997-4633.

May 1992

DIAC 92, Directions and Implications of Advanced Computing, May 2-3, Berkeley, Calif. Sponsor: Computer Professionals for Social Responsibility. Contact Douglas Schuler, Boeing Computer Services, MS 7L-64, PO 24346, Seattle, WA 98124-0346, home phone (206) 632-1659, office phone (206) 865-3832, fax (206) 865-2964, e-mail dschuler@june.cs.washington.edu.

Digital Equipment Computer Users Soc. Spring Seminars and Symp., May 2-8, Atlanta. Sponsor: US Chapter of DECUS. Contact DECUS, SHR 1-4/D33, 333 South St., Shrewsbury, MA 01545-4195, phone (508) 841-3341.

Third Rexx Symp. for Developers and Users, May 3-5, Annapolis, Md. Sponsor: Stanford Linear Accelerator Center. Contact Cathie Dager, SLAC, PO Box 4349, MS 97, Stanford, CA 94309, phone (415) 926-2904, fax (415) 926-3329, e-mail cathie@slacvm.slac.stanford.edu.

CICC 92, IEEE Custom Integrated Circuits Conf., May 3-6, Boston. Contact Roberta Kaspar, 1597 Ridge Rd. W, Suite 101C, Rochester, NY 14615, phone (716) 865-7164.

VHDL Int'l Users Forum Spring Conf., May 3-6, Scottsdale, Ariz. Contact Allen Dewey, IBM, PO Box 950, Dept. G75B, Bldg. 708, Poughkeepsie, NY 12602, phone (914) 435-6040, fax (914) 432-9407, e-mail adewey@pk705vmg.vnet.ibm.com.

CHI 92, Conf. on Human Factors in Computing, May 3-7, Monterey, Calif. Sponsor: ACM. Contact Assoc. for Computing Machinery, 11 W. 42nd St., New York, NY 10036, phone (212) 869-7440.

1992 IEEE Symp. on Research in Security and Privacy, May 4-6, Oakland, Calif. Sponsor: IEEE Computer Soc. Technical Committee on Security and Privacy. Contact Deborah Cooper, Unisys, 5731 Slauson Ave., Culver City, CA 90230, phone (310) 338-3727, e-mail cooper@culv.unisys.com.

ICOOMS 92, Int'l Conf. on Object-Oriented Manufacturing Systems, May 4-6, Calgary, Alta., Canada. Cosponsors: Soc. for Computer Simulation, Univ. of Calgary. Contact Douglas Norrie, Manufacturing Eng. Div., Eng. Faculty, Univ. of Calgary, 2500 University Dr. NW, Calgary, Alta., Canada T2N 1N4, phone (403) 220-5787, fax (403) 282-8406.

Second Workshop on Software Reliability, May 4-6, Ottawa, Canada. Cosponsors: IEEE Canadian Review et al. Contact Shirley Mills, Statistical Consulting Centre, Mathematics and Statistics Dept., Rm. 613 Dunton Tower, Carleton Univ., Ottawa, Ont., Canada K1S 5B6, phone (613) 788-3984, fax (613) 788-2148, e-mail stat_consulting@carleton.ca.

STOC 92, 24th Symp. the Theory of Computing, May 4-6, Victoria, B.C., Canada. Sponsor: ACM. Contact Mike Fellows, Computer Science Dept., Univ. of Victoria, Victoria, B.C., Canada V8W 2Y2, phone (604) 721-7299, e-mail mfellows@crs.uvic.ca.

Int'l Council for Computer Comm. Conf. on Intelligent Networks, May 4-6, Tampa, Fla. Contact Paul S. Bayless, AT&T Network Systems, Rm. 4M-612, 101 Crawfords Corner Rd., Holmdel, NJ 07733, phone (908) 949-4122, fax (908) 949-6222.

Food Processing Automation Conf., May 4-6, Lexington, Ky. Sponsor: Food and Process Eng. Inst. Contact Fred Payne, Univ. of Kentucky, 220 Agricultural Eng. Bldg., Lexington, KY 49546-0276, phone (606) 257-3000, ext. 220.

ECF 92, Eastern Communications Forum, May 4-6, Rye Brook, N.Y. Sponsor: Nat'l Eng. Consortium. Contact NEC, 303 E. Wacker Dr., Suite 740, Chicago, IL 60601, phone (312) 938-3500, fax (312) 938-8787.

STAR 92, Int'l Software Testing Analysis and Review Conf., May 4-7, Las Vegas, Nev. Cosponsors: Data Processing Management Assoc. et al. Contact Software Quality Eng., 3000-2 Hartley Rd., Jacksonville, FL 32257, phone (904) 268-8639, fax (904) 268-0733.

IEEE Infocom 92, 11th Conf. on Computer Comm., May 4-8, Florence, Italy. Cosponsor: IEEE Comm. Soc. Contact L. Fratta, Politecnico di Milano, c/o Cefriel, Via Emanueli, 15, 20126 Milano, Italy, phone 39 (2) 2399-3578, fax 39 (2) 2399-3587, e-mail fratta@imicefr.bitnet; or J. Kurose, Computer and Information Science Dept., Univ. of Massachusetts, Amherst, MA 01003, phone (413) 545-1585, e-mail kurose@cs.umass.edu.

Comp Euro 92, IEEE Int'l Conf. on Computer Systems and Software Eng., May 4-8, The Hague, The Netherlands. Cosponsors: IEEE Region 8, IEEE Benelux Section. Contact Patrick Dewilde, Delft Univ. of Tech., EE Dept., POB 5031, 2600 GA Delft, The Netherlands, fax 31 (15) 623-271; or G. Arink, Philips Medical Systems, PO Box 10000, 5680 DA Best, The Netherlands, fax 31 (40) 762-952.

NASA 92, Seventh Goddard Conf. on Space Applications of Artificial Intelligence, May 5-6, Greenbelt, Md. Sponsor: NASA Goddard Space Flight Center. Contact Jonathan Hartley, NASA GSFC, Code 522, Greenbelt, MD 20771, phone (301) 286-3150.

Nat'l Online Meeting and **Integrated Online Library Systems 92, May 5-7,** New York City. Contact Nat'l Online Meeting, Learned Information, 143 Old Marlton Pike, Medford, NJ 08055, phone (609) 654-6266, fax (609) 654-4309.

PTW 92, Second Pacific Test Workshop, May 5-8, Vancouver, B.C., Canada. Sponsor: IEEE Computer Soc. Technical Committee on Test Tech. Contact Andre Ivanov, Electrical Eng. Dept., Univ. of British Columbia, 2356 Mian Mall, Vancouver, B.C., Canada V6T 1Z4, phone (604) 822-6936, fax (604) 822-5949; or Mani Soma, phone (206) 685-3810, fax (206) 543-3842.

Data Processing Controllership Seminar, May 6-8, Las Vegas, Nev. Sponsor: Financial Management for Data Processing. Contact Terence A. Quinlan, FMDP, PO Box 27543, San Francisco, CA 94127, phone (415) 731-3706.

First Nat'l Conf. of the Computer Ethics Inst., May 7-8, Washington, DC. Held in cooperation with the IEEE Computer Soc. Committee on Public Policy. Contact CEI, 11 Dupont Cir. NW, Suite 900, Washington, DC 20036-1271.

1992 Am. Medical Informatics Assoc. Spring Congress, May 7-9, Portland, Ore. Contact AMIA, 4915 St. Elmo Ave., Suite 302, Bethesda, MD 20814.

Compotek 92, May 7-10. Contact Izfas, Sair Esref Belvari No. 50, Kültürpark, Izmir, Turkey, fax (51) 25-43-42.

CASE Data Interchange Format, May 8, New York City. Sponsor: CDIF Technical Committee of the Electronic Industries Assoc. Contact Patti Rusher, EIA, 2001 Pennsylvania Ave. NW, Washington, DC 20006, phone (202) 457-4962, fax (202) 457-4985.

IEEE Computer Soc. Workshop on Ethical Standards for the Profession, May 8-9, Washington, DC. Sponsor: IEEE Computer Soc. Committee on Public Policy. Contact IEEE Computer Soc., 1730 Massachusetts Ave. NW, Washington, DC 20036-1903, phone (202) 371-1013, fax (202) 728-9614; or Robert J. Melford, R.J. Melford and Associates, 24571 Tarazona, Mission Viejo, CA 92692-2350, phone (714) 830-1443, e-mail r.melford@ieee.org.

1992 IEEE Int'l Conf. on Robotics and Automation, May 10-15, Nice, France. Sponsor: IEEE Robotics and Automation Soc. Contact Harry Hayman, PO Box 3216, Silver Spring, MD 20918, phone (301) 236-5621, fax (301) 236-5621.

CLEO 92, 12th Conf. on Lasers and Electro-Optics and **QELS 92, Third Quantum Electronics and Laser Science Conf., May 10-15,** Anaheim, Calif. Cosponsors: IEEE

Lasers and Electro-Optics Soc. et al. Contact Optical Soc. of Am., 2010 Massachusetts Ave. NW, Washington, DC 20036, phone (202) 223-8130, fax (202) 223-1096.

ICSE 92, 14th Int'l Conf. on Software Eng., May 11-15, Melbourne, Australia. Cosponsors: IEEE Computer Soc. Technical Committee on Software Eng. et al. Contact IEEE Computer Soc., 1730 Massachusetts Ave. NW, Washington, DC 20036-1903, phone (202) 371-1013, fax (202) 728-0884; or A.Y. Montgomery, Computer Science Dept., Royal Melbourne Inst. of Tech., PO Box 2476V, Melbourne 3001, Victoria, Australia, phone 61 (3) 660-2943, fax 61 (3) 662-1617, e-mail aym@goanna.cs.rmit.oz.au.

Vision Interface 92 and **Graphics Interface 92, May 11-15**, Vancouver, B.C., Canada. VI 92 sponsor: Canadian Image Processing and Pattern Recognition Soc. GI 92 sponsor: Canadian Human-Computer Comm. Soc. Contact Kelly Booth or Alain Fournier, Computer Science Dept., Univ. of British Columbia, Vancouver, B.C., Canada V6T 1W5, phone (604) 822-8193 or 822-8990, fax (604) 228-5485, e-mail ksbooth@cs.ubc.ca or fournier@cs.ubc.ca.

Electro 92, May 12-14, Boston. Cosponsors: IEEE et al. Contact Electro 92, 8110 Airport Blvd., Los Angeles, CA 90045, phone (800) 877-2668.

Test and Design Expo, May 12-14, Somerset, N.J. Sponsor: Miller Freeman Expositions. Contact Carol Hurley, MFE, 1050 Commonwealth Ave., Boston, MA 02215-1135, phone (617) 232-3976, fax (617) 232-0854.

USING 92, May 12-14, Philadelphia. Sponsor: Unix Systems Information Network Group. Contact Don Ackerson, USING, PO Box 1077, Lisle, IL 60532, phone (708) 896-5777.

First European Conf. on Smart Structures and Materials, May 12-14, Glasgow, Scotland. Sponsor: Int'l Soc. for Optical Eng. Contact SPIE, PO Box 10, Bellingham, WA 98227-0010, phone (206) 676-3290, fax (206) 647-1445.

Ninth IEEE Workshop on Real-Time Operating Systems and Software, May 13-14, Pittsburgh. Cosponsors: IEEE Computer Soc. Technical Committee on Real-Time Systems, Office of Naval Research. Contact Marc Donner, IBM Research, PO Box 218, Yorktown Heights, NY 10598, phone (914) 784-7508 or (914) 945-2032, fax (914) 945-1234, e-mail donner@watson.ibm.com.

SID 92, Int'l Symp. and Seminar of the Soc. for Information Display, May 17-22, Boston. Contact Paul M. Alt, IBM T.J. Watson Research Center, PO Box 218, Yorktown Heights, NY 10598, phone (914) 945-2437, fax (914) 945-1974, e-mail alt@watson.ibm.com; Joyce Farrell, Hewlett-Packard Labs, PO Box 10490, Palo Alto, CA 94303-0969, phone (415) 857-2807, fax (415) 857-4320, e-mail farrell@hplabs.hp.com; or Jay Morreale, Palisades Inst. for Research Services, 201 Varick St., Suite 1140, New York, NY 10014, phone (212) 620-3371, fax (212) 620-3379.

Workshop on Interconnections within High-Speed Digital Systems, May 18-20, Santa Fe, N.M. Sponsor: IEEE Lasers and Electro-Optics Soc. Contact Ken Young, Bellcore, 445 South St., Rm. 2Q-150, Morristown, NY 07962-1910.

Third Eurographics Workshop on Rendering, May 18-20, Bristol, U.K. Sponsor: Univ. of Bristol. Contact Alan Chalmers, Computer Science Dept., Univ. of Bristol, University Walk, Bristol, BS8 1TR, UK, phone 44 (272) 303-109, fax 44 (272) 251-154, e-mail alan@compsci.bristol.ac.uk.

Introduction to Chargeback and DP Cost Accounting, May 18-20, Charleston, S.C. Sponsor: Financial Management for Data Processing. Contact Terence A. Quinlan, FMDP, PO Box 27543, San Francisco, CA 94127, phone (415) 731-3706.

Int'l Symp. on Optical Materials Tech. for Energy Efficiency and Solar Energy, May 18-22, Toulouse, France. Sponsor: Int'l Soc. for Optical Eng. Contact SPIE, PO Box 10, Bellingham, WA 98227-0010, phone (206) 676-3290, fax (206) 647-1445.

ISCA 92, 19th Int'l Symp. on Computer Architecture, May 19-21, Queensland, Australia. Cosponsors: IEEE, ACM SIGArch. Contact Jean-Luc Gaudiot, EE-Systems Dept., SAL 300, Univ. of Southern California, Los Angeles, CA 90089-0781, phone (213) 740-4484 or (213) 743-0249, fax (213) 740-4449, e-mail gaudiot@usc.edu or gaudiot@priam.usc.edu; or D. Abramson, CSIRO, 723 Swanston St., Melbourne 3000, Australia, phone 61 (3) 282-2666, e-mail david.abramson@mel.dit.csiro.au.

Sixth Int'l Conf. on Computer Science, May 20-22, Tunis, Tunisia. Sponsor: Assoc. Francaise pour la Cybernetique Economique et Technique. Contact Montasser Ouaili, Ecole Nationale des Sciences de l'Informatique, 16 Rue 8010 Quartier Montplaisir, 1002 Tunis Belvedere, Tunisia, phone 216 (1) 784-032, fax 216 (1) 787-827.

IEEE Int'l Symp. on Industrial Electronics, May 25-29, Xian, China. Cosponsors: IEEE Industrial Electronics Soc. et al. Contact Allen C. Chen, AT&T Bell Labs, Rm. 1E134A, Whippany Rd., PO Box 903, Whippany Road, NJ 07981-0903; J.D. Irwin, Electrical Eng. Dept., Auburn Univ., Auburn, AL 36849, phone (205) 844-1810, fax (205) 844-1809; or Sheng-wu Liu, Library, Northwestern Polytechnical Univ., Xian 710071, China, phone 1 (29) 533-71 ext. 2926, fax 1 (29) 711-959.

ISTCS, Israeli Symp. on the Theory of Computing and Systems, May 27-28, Haifa, Israel. Cosponsors: Assoc. of Technological Education in Electronics and Computer Science et al. Contact Michael Rodeh, Science and Tech., IBM Israel, Technion City, Haifa 32000, Israel, phone 972 (4) 296-205, fax 972 (4) 320-894.

MVL 92, 22nd Int'l Symp. on Multiple-Valued Logic, May 27-29, Sendai, Japan. Sponsors: IEEE Computer Soc. Technical Committee on Multiple-Valued Logic, Japan Research Group on Multiple-Valued Logic. Contact Tatsuo Higuchi, Electronic Eng. Dept., Tohoku Univ., Aoba, Aramaki, Sendai 980, Japan, phone 81 (022) 222-1800, fax 81 (022) 263-9406, e-mail thiguchi@higuchi.ecei.tohoku.ac.jp; or S.B. Silio, Electrical Eng. Dept., Univ. of Maryland, College Park, MD 20742, phone (301) 454-6839, fax (301) 454-1837, e-mail silio@eng.umd.edu.

Symp. on Assessment of Quality Software Development Tools, May 27-29, New Orleans. Sponsor: Tulane Univ. Contact Judy Lee, IBM, 1000 NW 51st. St., Boca Raton, FL 33432, phone (407) 982-1048; or Ez Nahouraii, IBM (798/089), 6321 San Ignacio Ave., San Jose, CA 95119, phone (408) 281-5741, e-mail eznah@stlvm7.iinus1.ibm.com.

ICCI, Fourth Int'l Conf. on Computing and Information, May 28-30, Toronto. Contact Peter E. Lauer, Computer Science and Systems, McMaster Univ., Hamilton, Ont., Canada, (416) 525-9140.

June 1992

FGCS 92, Fourth Int'l Conf. on Fifth-Generation Computer Systems, June 1-5, Tokyo. Cosponsors: Information Processing Soc. of Japan et al. Contact Hidehiko Tanaka, Electrical Eng. Dept., Univ. of Tokyo, 3-1 Hongo 7-chome, Bunkyo-ku, Tokyo 113, Japan, phone 81 (33) 3812-2111 ext. 6663, fax 81 (33) 3818-5706, e-mail tanaka@mtl.t.u-tokyo.ac.jp.

Sixth Israel Conf. on Computer Systems and Software Eng., June 2-3, Herzliya, Israel. Sponsors: IEEE Computer Soc. Israel Chapter, Information Processing Assoc. of Israel. Contact Conference Secretariat, Ortra, PO Box 50432, Tel Aviv 61500, Israel, phone 972 (3) 664-825, fax 972 (3) 660-952.

European Design for Testability Conf., June 2-4, Brugge, Belgium. Contact Thomas W. Williams, phone (303) 924-9912; or J. Jamison, phone 32 (2) 118-7511.

Euro ASIC 92, June 2-4, Paris. Contact Gabriele Saucier, Laboratoire CSI/Institut National, Polytechnic de Grenoble, 46 Ave. Felix Viallet/38000 Grenoble, France, phone (33) 76-574-687, fax (33) 76-503-421.

COGANN, Combination of Genetic Algorithms and Neural Networks Workshop, June 6, Baltimore. Sponsor: IEEE Neural

Networks Council. Contact Darrell Whitley, Computer Science Dept., Colorado State Univ., Fort Collins, CO 80523, phone (303) 491-5373, e-mail whitley@cs.colostate.edu.

Φ **STW 92, Southwest Test Workshop, June 7-9**, San Diego, Calif. Sponsor: IEEE Computer Soc. Technical Committee on Test Tech. Contact William R. Mann, Rockwell Int'l, MS 503-200, 4311 Jamboree, Newport Beach, CA 92660, phone (714) 833-4013, fax (714) 833-6680.

IJCNN 92, Int'l Joint Conf. on Neural Networks, June 7-11, Baltimore. Cosponsors: IEEE, Int'l Neural Network Soc. Contact IJCNN 92, 5665 Oberlin Dr. #110, San Diego, CA 92121, phone (619) 453-6222, fax (619) 535-3880.

Φ **DAC 92, 29th IEEE/ACM Design Automation Conf., June 8-12**, Anaheim, Calif. Contact DAC 92, MP Associates, 7490 Clubhouse Rd., Boulder, CO 80301, phone (303) 530-4333; or Dan Schweikert, Cadence Design Systems, 555 River Oaks Pkwy., Bldg. 4, San Jose, CA 95132, phone (408) 944-7297, e-mail dan@cadence.com.

Φ **Sixth Int'l Workshop on Scientific and Statistical Database Management, June 9-11**, Zurich, Switzerland. Cosponsor: ACM SIGMOD. Contact H. Hinterberger, Inst. for Scientific Computing, ETH Zentrum, CH-8093 Zurich, Switzerland, phone 41 (1) 254-7436, fax 41 (1) 262-3973.

Φ **IEA/AIE 92, Fifth Int'l Conf. on Industrial and Eng. Applications of Artificial Intelligence and Expert Systems, June 9-12**, Paderborn, Germany. Cosponsors: Univ. of Paderborn et al. Contact Moonis Ali, Univ. of Tennessee Space Inst., MS 15, Tullahoma, TN 37388-8897, phone (615) 455-0631, ext. 283, e-mail alif@udsivl.bitnet; or Fevzi Belli, Univ. of Paderborn, Postfash 1621, 4790 Paderborn, Germany, phone 49 (5251) 60-3282, fax 49 (5251) 602-519, e-mail belli@adt.uni-paderborn.de.

Φ **ICDCS 92, 12th Int'l Conf. on Distributed Computing Systems, June 9-12**, Yokohama, Japan. Cosponsor: Information Processing Soc. of Japan. Contact Ming T. (Mike) Liu, Computer and Information Science Dept., Ohio State Univ., 2036 Neil Ave., Columbus, OH 43210-1277, phone (614) 292-6552, fax (614) 292-9021, e-mail mike.liu@osu.edu; or Yutaka Matsushita, Instrumentation Eng. Dept., Keio Univ., 3-14-1 Hiyoshi, Kohoku-ku, Yokohama, Japan 223, fax 81 (045) 562-7625, e-mail on@inst.keio.ac.jp.

Φ **ITS 92, Second Int'l Conf. on Intelligent Tutoring Systems, June 10-12**, Montreal. Cosponsors: Univ. of Montreal et al. Contact Claude Frasson, Departement d'IRO, Universite de Montreal, CP 6128 Succ A. Montreal, Que., Canada H3C 3J7, phone (514) 343-7019, fax (514) 343-5834, e-mail frasson@iro.umontreal.ca.

Φ **CBMS 92, IEEE Symp. on Computer-Based Medical Systems, June 14-17**, Durham, N.C. Cosponsors: IEEE Computer Soc., IEEE Eastern North Carolina Section, Eng. in Medicine and Biology Soc. Contact James N. Brown, Jr., Research Triangle Inst., PO Box 12194, Research Triangle Park, NC 27709, phone (919) 541-6975, fax (919) 541-5945; or Peter Santago, Radiology Dept., Bowman Gray School of Medicine, Medical Center Blvd., Winston-Salem, NC 27157-1022, phone (919) 748-4260, fax (919) 748-2870, e-mail cbms@mrips.bgsm.wfu.edu.

Φ **NECC 92, 13th Nat'l Educational Computing Conf., June 15-17**, Dallas. Contact Cathy Norris, Computer Education and Cognitive Systems, Univ. of North Texas, PO Box 5155, UNT Station, Denton, TX 76203, phone (817) 565-2824, fax (817) 565-4425.

Φ **ICSI 92, Second Int'l Conf. on Systems Integration, June 15-18**, Morristown, N.J. Cosponsors: New Jersey Inst. of Tech. Computer and Information Science Dept., ACM et al. Contact Peter A. Ng, Computer and Information Science Dept., NJIT, University Heights, Newark, NJ 07102, phone (201) 596-3387, e-mail ng_p@vienna.njit.edu; or C.V. Ramamoorthy, Univ. of California, EECS Dept., Berkeley, CA 94720, phone (415) 642-4751, fax (415) 642-5775.

Φ **CVPR 92, IEEE Computer Society Conf. on Computer Vision and Pattern Recognition, June 15-18**, Champaign, Ill. Contact A. Rosenfeld, Center for Automation Research, Univ. of Maryland, College Park, MD 20742, e-mail ar@alv.umd.edu.

Φ **Parle 92, Parallel Architectures and Languages Europe, June 15-18**, Paris. Cosponsors: IEEE Region 8 et al. Contact Daniel Etiemble, LRI, Bat 490, Univ. of Paris Sud, 91405 Orsay Cedex, France, fax 33 (1) 6941-6586, e-mail de@lri.lri.fr.

Compass 92, Seventh Conf. on Computer Assurance, June 15-18, Gaithersburg, Md. Cosponsors: IEEE, IEEE Aerospace and Electronic Systems Soc. Contact Robert Ayers, ARINC Research, 2551 Riva Rd., Annapolis, MD 21401, phone (301) 266-4741, fax (301) 266-4040.

Φ **Computer Security Foundations Workshop V, June 16-18**, Franconia, N.H. Contact Leonard LaPadula, Information Security Technical Center, Mitre Corp., Bedford, MA 01730-0208, phone (617) 271-3261, e-mail ljl@mitre.org; or Ravi Sandhu, ISSE Dept., George Mason Univ., Fairfax, VA 22030-4444, phone (703) 993-1659, e-mail sandhu@sitevax.gmu.edu.

Φ **SEKE 92, Fourth Int'l Conf. on Software Eng. and Knowledge Eng., June 17-19**, Capri, Italy. Cosponsors: Univ. of Salerno et al. Contact Shi-Kuo Chang, Computer Science Dept., Univ. of Pittsburgh, Pittsburgh, PA 15260, phone (412) 624-8493; or Genoveffa Tortora, Dipartimento di Informatica ed Applicazioni, Univ. di Salerno, I-84081 Baronissi, Salerno, Italy, phone 39 (89) 822-333, fax 39 (81) 482-745, e-mail jentor@udsab.dia.unisa.it or usditort@inacriai.bitnet.

Φ **Int'l Workshop on Hardware Fault-Tolerance in Multiprocessors, June 22-23**, Amherst, Mass. Sponsor: IEEE Computer Soc. Technical Committee on Fault-Tolerant Computing. Contact Donald S. Fussell, Computer Science Dept., TAY 2.124, Univ. of Texas at Austin, Austin, TX 78712-1188, phone (512) 471-9719, fax (512) 471-7028, e-mail fussell@cs.utexas.edu.

Φ **Seventh IEEE Conf. on Structure in Complexity Theory, June 22-25**, Boston. Sponsor: IEEE Computer Soc. Technical Committee for Math. Foundations of Computing. Contact J. Royer, School of Computer and Information Science, Syracuse Univ., Syracuse, NY 13244, fax (315) 443-1954, e-mail royer@top.cis.syr.edu.

Φ **LICS, Seventh IEEE Symp. on Logic in Computer Science, June 22-25**, Santa Cruz, Calif. Sponsor: IEEE Computer Soc. Technical Committee on Mathematical Foundations in Computing. Contact Robert L. Constable, Computer Science Dept., Cornell Univ., 4149 Upson, Ithaca, NY 14853, phone (607) 255-9204, fax (607) 255-4428.

Φ **CGI 92, Computer Graphics Int'l Conf., June 22-26**, Tokyo. Cosponsors: Computer Graphics Soc. et al. Contact R.A. Earnshaw, Computer Graphics, Univ. of Leeds, Leeds LS2 9JT, England, phone 44 (532) 335-414, fax 44 (532) 336-017; or Toki Takahashi, Autonomous Robot Systems Lab, NTT Human Interface Labs, 3-9-11 Midori-cho, Musashino-shi, Tokyo 180, Japan, fax 81 (422) 59-2245, e-mail cgi92@nttarm.ntt.jp (Internet).

Φ **RSP, Third Int'l Workshop on Rapid System Prototyping, June 23-25**, Research Triangle Park, N.C. Cosponsors: IEEE Computer Soc. Technical Committees on Design Automation, Simulation, and Test Tech. Contact Nick Kanopoulos, Center for Systems Eng., Research Triangle Inst., 3040 Cornwallis Rd., Research Triangle Park, NC 27709, phone (919) 541-7341, fax (919) 541-6515, e-mail rsp@rti.rti.org.

Φ **10th Software Reliability Symp., June 25-26**, Denver. Sponsor: IEEE Reliability Soc. Denver Chapter. Contact Rich Karcich, Storage Tech. Corp., 2270 S. 88th St., MS 2286, Louisville, CO 80028-5319, phone (303) 673-6223, fax (303) 673-3353.

IEEE Intelligent Vehicles 92, June 29-July 1, Ypsilanti, Mich. Sponsor: IEEE Industrial Electronics Soc. Contact Ichiro Masaki, Computer Science Dept., GM Research Labs, 30500 Mound Rd., Warren, MI 48090-9055, phone (313) 986-1466, fax (313) 986-9356, e-mail masaki@gmr.com.

July 1992

Φ **1992 Workshop on Fault-Tolerant Parallel and Distributed Systems, July 6-7,** Amherst, Mass. Cosponsor: Univ. of Massachusetts at Amherst. Contact Donald S. Fussell, Computer Science Dept., Univ. of Texas at Austin, TAY 2.124, Austin, TX 78712-1188, phone (512) 471-9719, fax (512) 471-7028, e-mail fussell@cs.utexas.edu.

Φ **CASE 92, Fifth Int'l Workshop on Computer-Aided Software Eng., July 6-10,** Montreal. Contact Gene Forte, CASE Outlook, 11830 Kerr Pkwy. #315, Lake Oswego, OR 97035, phone (503) 245-6880, fax (503) 245-6935, e-mail g.forte@compmail.com; Nazim Madhavji, School of Computer Science, McGill Univ., 3480 University St., Montreal, Que., Canada H3A 2A7, phone (514) 398-3740, fax (514) 398-3883, e-mail case@softeng.cs.mcgill.ca; or John Jenkins, City Univ. of London, Business Systems Analysis Dept., London, UK, fax 44 (71) 608-1270.

Iros 92, Fifth Int'l Conf. on Intelligent Robots and Systems, July 7-10, Raleigh, N.C. Cosponsors: IEEE Industrial Electronics Soc., Robotics Soc. of Japan et al. Contact Avi Kak, School of Electrical Eng., Purdue Univ., West Lafayette, IN 47907, phone (317) 494-3551, fax (317) 494-6440, e-mail kak@ecn.purdue.edu; or Kazuo Tanie, Mechanical Eng. Lab, MITI, 1-2 Namiki, Tsukuba-shi, Ibaraki-ken, 305, Japan, phone 81 (298) 54-2656, fax 81 (298) 54-2518, e-mail m1750@mel.go.jp.

Φ **FTCS 22, 22nd Int'l Symp. on Fault-Tolerant Computing, July 8-10,** Boston. Cosponsor: Univ. of Massachusetts at Amherst. Contact Dhiraj K. Pradhan, Electrical and Computer Eng. Dept., Univ. of Massachusetts at Amherst, Amherst, MA 01003, phone (413) 545-0160, fax (413) 545-4611, e-mail pradhan@ecs.umass.edu.

Logic at Tver 92: Logical Foundation of Computer Science, July 20-24, Tver, Russia. Cosponsors: IEEE et al. Contact MA Taitslyn, Univ. of Tver, 33 Zhelyahova St., Tver 170013, Russia.

Φ **Siggraph 92, 19th Int'l Conf. on Computer Graphics and Interactive Techniques, July 26-31,** Chicago. Sponsor: ACM. Contact Siggraph 92, 401 N. Michigan Ave., Chicago, IL 60611, phone (312) 321-6830, fax (312) 321-6876, e-mail info92@siggraph.org.

August 1992

Φ **17th Congress Workshop, Aug. 2-14,** Washington, DC. Sponsor: United Nations. Contact Lawrence W. Fritz, GE Aerospace, PO Box 8048-10A26, Philadelphia, PA 19101, phone (215) 531-3205, fax (215) 962-3698.

Φ **ASAP 92, Int'l Conf. on Application-Specific Array Processors, Aug. 4-7,** Berkeley, Calif. Sponsor: Univ. of California at Berkeley. Contact Jose Fortes, School of Electrical Eng., Purdue Univ., West Lafayette, IN 47907, phone (317) 494-3646, fax (317) 494-6440, e-mail fortes@ecn.purdue.edu; or Mary Stewart, EECS Dept., Cory Hall, Univ. of California, Berkeley, CA 94720.

Φ **Hot Chips IV, Symp. on High-Performance Chips, Aug. 9-11,** Stanford, Calif. Sponsor: IEEE Computer Soc. Technical Committee on Microprocessors and Microcomputers. Contact Glen Langdon, phone (408) 459-2212, e-mail langdon@cse.ucsc.edu.

Φ **Crypto 92, Aug. 16-20,** Santa Barbara, Calif. Sponsor: Int'l Assoc. of Cryptologic Research. Contact Spyros S. Magliveras, Computer Science and Eng. Dept., Univ. of Nebraska, Lincoln, NE 68588-0115, phone (402) 472-5005, fax (402) 472-7767, e-mail spyros@helios.unl.edu.

ICPP 92, 21st Int'l Conf. on Parallel Processing, Aug. 17-21, St. Charles, Ill. Sponsor: Pennsylvania State Univ. Contact Tse-yun Feng, Penn State Univ., ECE Dept., EE East Bldg., University Park, PA 16802, phone (814) 863-1469, e-mail tfeng@ccl.psu.edu.

Φ **Int'l Workshop on Distributed Data Management, Aug. 19-21,** Edmonton, Alta., Canada. Cosponsors: Canadian Information Processing Soc. et al. Contact M. Tamer Ozsu, Computing Science Dept., Univ. of Alberta, Edmonton, Alta., Canada T6G 2H1, phone (403) 492-1071, fax (403) 492-1071.

Φ **VLDB 92, 18th Int'l Conf. on Very Large Data Bases, Aug. 23-27,** Vancouver, B.C., Canada. Cosponsors: VLDB Endowment, Canadian Information Processing Soc. et al. Contact Paul Sorenson, Computer Science Dept., Univ. of Alberta, Edmonton, Alta., Canada T6G 2H1, phone (403) 492-4589, fax (403) 492-1071, e-mail vldb92@cs.ualberta.ca or sorenson@cs.ualberta.ca.

Φ **Int'l Workshop on Hardware-Software Codesign, Aug. 31-Sept. 2,** Estes Park, Colo. Cosponsors: ACM et al. Contact Joanne E. DeGroat, Ohio State Univ., 205 Neil Ave., Columbus, OH 43210, phone (614) 292-2439, e-mail degroat@ee.eng.ohio-state.edu.

September 1992

Φ **Euro DAC 92, European Design Automation Conf., Sept. 7-10,** Hamburg, Germany. Cosponsors: IEEE Circuit and Systems Soc. et al. Contact A. Zylbersztejn, Bull SA, 121 Ave. De Malakoff, 75116 Paris, France, phone 33 (14) 502-9583, fax 33 (14) 502-9307.

Congress 92, 12th World Computer Congress, Sept. 7-11, Madrid. Sponsor: Int'l Federation for Information Processing. Contact Grupo Geyseco, Congress 92, Mauricio Legendre, 4-8.° G., E-28046 Madrid, Spain, fax 34 (1) 323-4936, e-mail ifip92@dit.upm.es.

ICIP 92, Second Singapore Int'l Conf. on Image Processing, Sept. 7-11, Singapore. Cosponsors: IEEE Singapore Section et al. Contact Technical Programme Chair, ICIP 92, IEEE Singapore Section, 200 Jalan Sultan, #11-03 Textile Centre, Singapore 0719, phone (65) 291-9690, fax (65) 292-8596.

Φ **HPDC I, First Int'l Symp. on High-Performance Distributed Computing, Sept. 9-10,** Syracuse, N.Y. Cosponsor: Syracuse Univ. Contact Geoffrey C. Fox, NPAC, Syracuse Univ., Rm. 3-131, 111 College Pl., Syracuse, NY 13244-4100, phone (315) 443-4741, fax (315) 443-1973, e-mail hpdc@nova.npac.syr.edu.

Φ **LCN 92, 17th Conf. on Local Computer Networks, Sept. 13-16,** Minneapolis. Sponsor: IEEE Computer Soc. Technical Committee on Computer Comm. Contact Steve Bell, Hughes LAN Systems, MS 392, 1072 S. Saratoga Svale Rd., San Jose, CA 95129, phone (415) 966-7926, fax (415) 966-7990, e-mail sbell@hls.com.

Fifth Digital Signal Processing Workshop, Sept. 13-16, Starved Rock State Park, Ill. Sponsor: IEEE Signal Processing Soc. Contact Mark J.T. Smith, School of Electrical Eng., Georgia Inst. of Tech., Atlanta, GA 30332-0250, phone (404) 894-6291.

Φ **SESS, Software Eng. Standards Symp., Sept. 13-17,** Brighton, UK. Contact Takis Katsoulakos, Lloyd's Register, Lloyd's Register House, 29 Wellesley Rd., Croydon, CR0 2AJ, UK, phone (081) 681-4774.

Φ **DCCA 3, Third Working Conf. on Dependable Computing for Critical Applications, Sept. 14-16,** Mondello, Italy. Sponsor: Int'l Federation for Information Processing Working Group 10.4. Contact Carl E. Landwehr, Code 5540, Naval Research Lab, Washington, DC 20375-5000, phone (202) 767-3381, fax (202) 404-7942; or Luca Simoncini, Dipartimento di Ingegneria dell'Informazione, Univ. of Pisa, Via Diotisalvi 2, 56100 Pisa, Italy, phone 39 (50) 593-443 or 550-100, fax 39 (50) 554-342, e-mail simon@icnucevm.cnuce.cnr.it.

Φ **ICARCV 92, Second Int'l Conf. on Automation, Robotics, and Computer Vision, Sept. 15-18,** Singapore. Cosponsors: Singapore Institution of Engineers et al. Contact ICARCV Secretariat, Associated Conventions and Exhibitions, 204 Bukit Timah Rd., #04-00, Boon Liew Bldg., Singapore 0922, phone (65) 799-5470, fax (65) 791-2687, e-mail emital@ntuvax.bitnet.

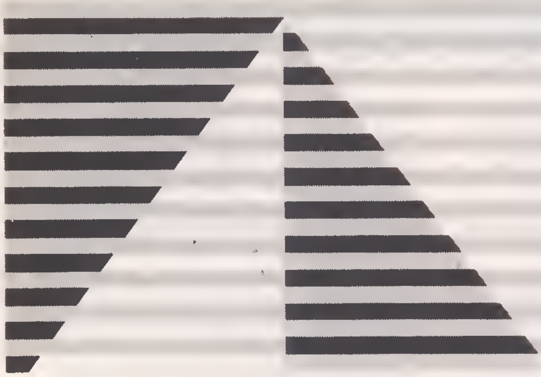

THE 19TH ANNUAL INTERNATIONAL SYMPOSIUM ON COMPUTER ARCHITECTURE

Gold Coast, Australia
May 18-22 1992

sponsored by

Association for Computing Machinery/ SIGARCH

IEEE Computer Society/IEEE TCCA

Institute of Electrical and Electronics Engineers (IEEE)

REGISTRATION FORM

First Name ______ Last Name ______
Company/Institution ______
Address ______
City/State/Zip/Ctry ______
Tel. ______ E-mail ______
Badge ______ Dietary or Special Needs ______
Check appropriate boxes to indicate member status
ACM ☐ IEEE-CS ☐
Member Number ______

Fees ($US)	Prior to April 27			After April 27		
	Student	Member	Nonmember	Student	Member	Nonmember
Conference Reg.	150	340	425	200	410	510
Full day tutorials		320	400		385	480
Half day tutorials		200	250		240	300

Tutorials	Full Day	Morning	Afternoon
Monday	1	2 or 3	4 or 5
Friday	6	7	8

Conference registration $ ______
Monday Tutorial number ______ fee $ ______
Tutorial number ______ fee $ ______
Friday Tutorial number ______ fee $ ______
Tutorial number ______ fee $ ______
Additional Dinner tickets # ______ × $50 = $ ______
TOTAL FEE ENCLOSED US $ ______
Payment method:
Credit Card ______ Check ______ Bank draft ______
For Credit Card Payment:
Please debit my credit card for the amount of US$ ______
Credit Card: VISA | American Express | Mastercard
Card Holder ______
Card Number ______ Expiry Date ______
Signature ______ Date ______

FOR PAYMENT IN U.S. FUNDS ONLY

Make checks payable to ACM/IEEE-CS 19TH ISCA and forward with completed registration form to:

ACM/IEEE-CS 19TH ISCA
c/o John Feo, L-306
Lawrence Livermore National Laboratory
P.O. Box 808,
Livermore, CA 94550,
USA
email: feo@sisal.llnl.gov

FOR PAYMENT IN AUSTRALIAN DOLLARS ONLY

TOTAL FEE US$ ______ x 1.47 = A$ ______

Complete the above conversion calculation, make checks for the converted amount payable to ACM/IEEE-CS 19TH ISCA and forward with completed registration form to:

ACM/IEEE-CS 19TH ISCA
c/o Rhys Francis
Div. of Information Tech., CSIRO
723 Swanston St.,
Carlton, Victoria, 3053,
Australia
email: rhys@mel.dit.csiro.au

Cancellations must be made in writing and received by April 27, 1992. Registration includes access to all technical sessions, one copy of the proceedings, coffee breaks and social functions. Proof of full time student enrollment must accompany registration form. Student registration does not include social functions. Tutorial registration includes copies of notes and coffee breaks. Tutorials may be cancelled due to lack of registration.

For further information:

*Conference—*Jean-Luc Gaudiot at (213) 740-4484, e-mail: isca92@usc.edu
*Workshops—*Lubomir Bic at (714) 856-5248, e-mail: bic@ics.uci.edu
Hotel— University Park Hotel at +61-75-955-000•FAX +61-75-525-091
Tutorials— Edward Lee at (510) 642-0455 e-mail: eal@joel.berkeley.edu.

CAREER OPPORTUNITIES

RATES: $15.00 per line (ten lines minimum). Average five typeset words per line, eight lines per column inch. Add $10 for box number. Send copy at least one month prior to publication date to: **Marian B. Tibayan, Classified Advertising, COMPUTER Magazine, 10662 Los Vaqueros Circle, PO Box 3014, Los Alamitos, CA 90720-1264; (714) 821-8380; fax: (714) 821-4010.**

In order to conform to the Age Discrimination in Employment Act and to discourage age discrimination, *COMPUTER* may reject any advertisement containing any of these phrases or similar ones: "...recent college grads...," "...1-4 years maximum experience...," "...up to 5 years experience," or "...10 years maximum experience." *COMPUTER* reserves the right to append to any advertisement without specific notice to the advertiser. Experience ranges are suggested minimum requirements, not maximums. *COMPUTER* assumes that since advertisers have been notified of this policy in advance, they agree that any experience requirements, whether stated as ranges or otherwise, will be construed by the reader as minimum requirements only. *COMPUTER* encourages employers to offer salaries that are competitive, but occasionally a salary may be offered that is significantly below currently acceptable levels. In such cases the reader may wish to inquire of the employer whether extenuating circumstances apply.

UNIVERSITY OF OREGON
Senior Faculty Position/ Department Head

The Department of Computer and Information Science invites applications for a senior faculty position created by a new state Centers of Excellence award. We are seeking a person who will be an active leader in the department, willing to serve one or more terms as department head and also play a key role in relations to the computer industry. Applicants should have a Ph.D. in computer science or related field and a distinguished record of teaching and research in the area of parallel processing (including parallel architectures, languages, and performance modeling) or human-computer interaction (including computer graphics and scientific visualization). Our department has 14 other faculty positions (including one other new position for which we are currently recruiting), approximately 20 Ph.D. students, 50 M.S. students, and 150 B.S. students. We have strong research programs in parallel and distributed systems, computer graphics, user interfaces, programming languages, software engineering, artificial intelligence, and theoretical computer science, and active interdisciplinary ties with other on-campus groups in the fields of cognitive science, neuroscience, economics, biology, physics, and mathematics. We offer a modern computing environment (a MasPar MP-1100, two Sequent Symmetry multiprocessors, and dozens of Sun and HP workstations) housed in a new computer science building.

Review of applications will begin on April 15 and continue until the position is filled. The position is available September, 1992, with a target date to fill the position by January, 1993. Qualified applicants should send their curriculum vitae and the names of at least three references to: Professor John S. Conery, Faculty Search Committee, Department of Computer and Information Science, University of Oregon, Eugene, OR 97403-1202. For more information send e-mail to conery@cs.uoregon.edu or phone (503) 346-3973. We especially encourage applications from women and minorities; the University of Oregon is an Equal Opportunity/Affirmative Action Employer committed to cultural diversity.

ACADEMIA SINICA
TAIWAN, REPUBLIC OF CHINA
Institute of Information Science

Applications are invited for research position in Institute of Information Science, Academia Sinica, Ph.D. in Computer Science or closely related fields required. Demonstratable research ability necessary. Applicants for senior positions must have proven research record. All fields in Computer Science are welcome.

The Institute offers a good research environment. No duty of teaching. Facilities include a 32-node NCUBE,2 parallel supercomputer, many SUN, SGI, and E&S workstations. An easily accessible ETA-10Q supercomputer is in the Academia Sinica.

Interested people please send application to Dr. Lin-Shan Lee, Director, Institute of Information Science, Academia Sinica, Taipei, Taiwan, 11529, Republic of China. Fax: (011-886-2) 782-4814.

NAVAL POSTGRADUATE SCHOOL
Chair of Department

The Computer Science Department of the Naval Postgraduate School seeks a Departmental Chair, with term beginning no later than 30 September 1992. The ideal candidate holds a Ph.D. in Computer Science, is from an academic environment, and maintains a continuing, strong involvement in research. We desire a candidate who facilitates the activities of departmental faculty and exerts leadership in developing department and interdisciplinary programs. The candidate should possess strong interpersonal and written communication skills. The candidate must be an excellent teacher.

The department offers graduate degrees with specialization tracks in artificial intelligence and robotics, computer graphics and visual simulation, computer systems and architectures, database and data engineering, and software engineering. The department currently has major applied research funding in software engineering, computer graphics/virtual reality, and artificial intelligence/robotics. We believe that our future departmental growth areas will be in the areas of software engineering/real-time systems, computer graphics/virtual reality, and secure distributed systems. Consequently, we prefer a chair whose research area matches with one of our areas of growth though we will consider exceptional candidates from any area. U.S. citizenship is required. Applicants interested in the position should forward their resume to:

Associate Professor Michael J. Zyda
Chair, Chair Search Committee
Naval Postgraduate School
Code CS/Zk, Department of Computer Science
Monterey, California 93943-5100
(408) 646-2305
Email: zyda@trouble.cs.nps.navy.mil

The Naval Postgraduate School is an Equal Opportunity/Affirmative Action Employer.

SENIOR SYSTEMS CONSULTANT

40 hours per week.; 8:00 a.m. - 5:00 p.m.; 10 hours overtime per week; Salary: $32.50 per hour; Overtime; $32.50 per hour.

Responsible for the design and implementation of computer based systems; designing and programming distributed database systems for business and engineering applications which can be integrated across a variety of communication networks; consulting with clients to recommend appropriate technology; marketing and promoting programs and services; maintaining, upgrading and redesigning computer systems; interviewing potential candidates for employment; supervising professional consultants' work; training professionals and personnel on the use of various computer systems and programs; serving as a liaison between companies; collecting data and preparing feasibility studies. Requirements: Bachelors Degree in Computer Science and at least one year of experience in job offered or one year of experience as a Systems Analyst/Consultant. Of the experience required applicant must have at least six months experience training at least three professionals in computer program use. Required experience must include at least one year of system design and construction experience with five successful system installations. Of coursework required applicant must have taken and passed at least one undergraduate course in database design and one course in structured systems analysis. Undergraduate transcript required. Apply at the Texas Employment Commission, Dallas, Texas, or send resume to the Texas Employment Commission, TEC Building, Austin, Texas 78778, J.O. #6687519. Ad Paid by an Equal Employment Opportunity Employer.

THE UNIVERSITY OF MINNESOTA MORRIS
Division of Science and Mathematics

The University of Minnesota, Morris, Division of Science and Mathematics seeks applications for a temporary one-year position, to teach undergraduate courses in computer science and mathematics starting September 16, 1992. The typical teaching load is two courses per quarter. Participation in division and campus governance is also expected. Appointments will be at the Assistant Professor level for those having the Ph.D. and at the instructor level for others. Salary is negotiable and depends upon qualifications. Reappointment for additional years is possible, depending on the availability of funds and the individual's performance. Minimum qualifications are a M.S. or M.A. degree in computer science or mathematics (Ph.D. preferred) with ability to teach in both areas and at least one year of teaching experience (graduate TA experience acceptable). Preference will be given to candidates who demonstrate potential for excellence in teaching. Send letter of application, resume, transcripts and three letters of recommendation via U.S. mail by May 1, 1992 to Dr. Michael Korth, Chair, Division of Science and Mathematics, University of Minnesota, Morris, MN 56267-2128. Inquiries can also be sent to A. Lopez, search committee chair, at andy@caa.mrs.umn.edu or andy@umnmor. The University of Minnesota is an equal opportunity educator and employer.

PROGRAMMER / ANALYST

Design, develop, upgrade, implement, and maintain computer applications software using AS/400 and RPG/400. Analyze, design, and develop computer systems. Interface with clients managerial and technical people to understand problems and requirements. Prepare technical reports, memoranda, and manuals. Must be able to use IBM System 36 and RPG II. BS in math or comp. science. 2 years experience in job offered using IBM System 36 and RPG II, AS/400, and RPG/400. Must be willing to accept temporary assignments anywhere and anytime as may be required by clients. 40 hrs./wk. 8 to 5. $519.23/wk. Job location: Maitland, Florida. Contact or send resume to Orlando Job Service Office, 3421 Lawton Road, Orlando, Florida 32803. Job Order #FL 0561915.

THE SWISS FEDERAL INSTITUTE OF TECHNOLOGY
Two Junior Faculty Positions in Computer Science (Informatik)

The Swiss Federal Institute of Technology in Zurich (ETH Zurich) invites applications for two Junior Faculty Positions in Computer Science (Informatik). One of the positions is designated for the area of Digital Circuits and Computer Architecture. The field of research and teaching of the other position is open, although preference will be given to software and application oriented areas, for example computer graphics, parallelism, or real-time information systems.

These positions offer the opportunity for independent research and teaching in a stimulating scientific environment. They are challenging posts for young scientists at the assistant professor level, available for three years in the first instance and renewable for another three years. Funding for assistant positions and equipment is available.

Applicants should have a Ph.D. in Computer Science or a related area, a record of independent and successful research, possibly teaching experience and the readiness to participate in the teaching of Computer Science at all university levels. The official language of ETH is German, although new professors may teach courses in English for a limited time.

Applications with curriculum vitae, references, and a short description of research interests should be sent no later than **April 30, 1992**, to the President of ETH Zurich, Prof. Dr. J. Nuesch, ETH Zentrum, CH-8092 Zurich. In view of the current composition of its faculty, ETH Zurich specifically encourages female candidates to apply.

Additional information about the positions can be obtained from the chairman of the Computer Science Department, Prof. Dr. C.A. Zehnder, Departement Informatik, ETH Zentrum, CH-8092 Zurich, Switzerland (zehnder@inf.ethz.ch).

THE UNIVERSITY OF ALABAMA

Applications are invited for a tenure track position in the Department of Computer Science at the Assistant Professor level to begin August 16, 1992. The University of Alabama Computer Science Department offers degrees at all levels. Applicants should have broad teaching and research interests within one or more areas of computer science. A Ph.D. in Computer Science is preferred, but applications will be accepted from those with the Ph.D. in a related field and extensive experience in Computer Science. Salary is commensurate with credentials. Applicants should submit a resume and the names of three references to Dr. Hui-Chuan Chen, Department of Computer Science, Box 870290, Tuscaloosa, AL 35487-0290 (Internet inquiries: chen@cs.ua.edu). Review of applications will begin on April 25, 1992 and will continue until the post is filled. The University of Alabama is an equal opportunity/affirmative action employer.

FACILITIES DIRECTOR

The Department of Computer Science is seeking to hire a Facilities Director for the departmental research and instructional facility (capitalization about $9M). The main responsibility is to coordinate with the department chairman and faculty to develop the facility and to provide an outstanding research and educational facility. Specific responsibilities include:

- planning and managing an annual budget of approximately $830,000.
- supervising the hardware, software, and operations managers of the facility.
- planning, specifying, procuring, replacing, and installing computer facilities, systems, and networks.
- negotiating and purchasing services and supplies contracts, as well as developing close interactions with amenable vendors on state-of-the-art products.
- working with the College of Engineering, Computer Center, other colleges and the University administration with regard to general computing issues.

A Ph.D. degree in Computer Science or a closely-related discipline is preferred, as is a minimum of five years experience in running a large facility. A demonstrated interest in helping map curricula onto workstations is also preferred.

Please send curriculum vita or resume and names of three references to: Shawn Darby, Administrative Manager, Department of Computer Science, 3190 Merrill Engineering Building, University of Utah, Salt Lake City, Utah 84112. The University of Utah is an Equal Opportunity, Affirmative Action Employer and encourages nomination and applications from women and minorities.

OREGON STATE UNIVERSITY
The Department of Electrical and Computer Engineering

The Department of Electrical and Computer Engineering invites applications for the following tenure-track faculty position in computer engineering. The opening is immediate. Associate or Full Professor applicants require a distinguished teaching and research record appropriate for the title. Candidates should have an earned doctorate in Electrical or Computer Engineering and are expected to have a strong commitment to high quality undergraduate and graduate teaching and to the development of a sponsored research program.

Areas of research activity include: computer architecture, parallel processing, VLSI array processing, performance modeling and analysis, data communication and computer networks.

With a faculty of 25, the Department of ECE enrolls about 425 undergraduate and 120 MS/Ph.D. students. The department offers an ABET accredited program in computer engineering. High technology corporations such as Hewlett-Packard, Intel, Mentor Graphics, and Tektronix have major operation in the area and provide support for the electrical and computer engineering programs. The Department has modern facilities housed in a new building. Located in the Willamette Valley 80 miles south of Portland, OSU and the city of Corvallis offer a beautiful and unspoiled environment and many cultural activities.

Applications must include a comprehensive resume, a list of three to five professional references, and a letter of interest. Your letter must indicate clearly the position for which you are applying. Please send material to Chairman, ECE Search Committee, ECE Department, Oregon State University, Corvallis, OR 97331-3211. Review will begin immediately and continue until the positions are filled.

Oregon State University is an Affirmative Action/Equal Opportunity Employer and complies with Section 504 of the Rehabilitation Act of 1973. OSU has a policy of being responsive to the needs of dual career couples.

UNIVERSITY OF CALIFORNIA, DAVIS

Faculty Positions in Electrical and Computer Engineering

The Department of Electrical and Computer Engineering at UC Davis invites applications for a tenure track faculty position. The area of interest is Computer Engineering with emphasis in advanced hardware and digital design topics such as logic synthesis, rapid prototyping, application specific processor design, timing verification, and programmable logic devices. The position is at the assistant professor level.

The department, with 33 faculty members and 140 full-time graduate students, is experiencing rapid growth. We have exceptional teaching and research facilities and are preparing to move into a large, new building which is currently under construction. Salary and benefits are extremely attractive.

Davis is a pleasant, family-oriented community near Sacramento, within easy driving distance to Silicon Valley, the Lawrence Livermore National Laboratory, San Francisco, the Pacific Ocean, and the Sierra Nevada Mountains.

We are seeking individuals with strong records in research, with ambitious plans and evidence of great promise. All faculty are expected to have a strong commitment to teaching at all degree levels, and to demonstrate the ability to attract significant research support.

The positions require a Ph.D. degree or equivalent, and are open until filled; but in order to assure consideration, applications should be received by May 1, 1992. Send a resume and the names of at least three references to:

Professor S. Louis Hakimi, Chair
Attention: Faculty Search Committee
Department of Electrical Engineering and Computer Science
University of California
Davis, CA 95616

The University of California, Davis, is an equal opportunity/affirmative action employer.

PRODUCT DEVELOPMENT ENGINEER I

Definition and design of VLSI (Very Large Scale Integration) 32 bit microprocessor architectures. This includes logic design, simulation and verification and circuit design, sizing (utilizing HSPICE) verification, and debug via hardware and software. Participates in layout verification, schematic capture and subsequent data bases. Support product debug and characterization to determine product performance, process and test sensitivities. Utilizes the C programming language on a Unix operating system. Provides central engineering support for the product(s) through all phases of development and/or manufacturing. Must have MS in Computer or Electrical Engineering. Must have successfully completed coursework in the following areas: VLSI Design; Circuit Design; Circuit Simulation; and Computer Architectures. Knowledge of the C programming language and Unix operating system required. 40 hours per week, 8:00 a.m. to 5:00 p.m., $42,180 per year. Apply at the Texas Employment Commission, Austin, Texas, or send resume to the Texas Employment Commission, TEC Building, Austin, Texas 78778, J.O. #6687138. Ad Paid by an Equal Employment Opportunity Employer.

COMPUTER SYSTEMS ADMINISTRATOR

Data processing consulting firm seeks Computer Systems Administrator to work in Denver, CO area to perform all aspects of computer network systems administration and applications, including network configuration, programming, and trouble-shooting, in a telecommunications environment. Includes hardware and software configuration, design and programming, including working with Integrated Service Digital Network (ISDN); distributed systems, such as remote file sharing, file transfer, and distributive processing; Oracle relational databases; Conversant Voice Information System; and X-window systems in a UNIX System V environment, using AT&T 3Bs, PC 386 hardware, and Sun workstations and Sun servers. Use a variety of computer languages, including C, Pro C, and UNIX shell scripts. Requires M.S. in Computer Science and one year experience in computer network and systems administration in a telecommunications environment, including the use of ISDN, distributed systems, UNIX System V and Oracle relational database. $37,500/year; 8:00am-5:00pm, M-F. Respond by resume no later than May 1, 1992 to Colorado Department of Labor & Employment, Division of Employment & Training, 600 Grant, Suite 900, Denver, CO 80203, ATT: James Shimada, and refer to Job Order No. CO3773915.

COLUMBIA UNIVERSITY

Department of Computer Science

We are anticipating the possibility of one or more tenure track openings. We invite applications from exceptional candidates at the assistant professor level. While we will consider applicants in all areas, we are particularly interested in hardware, including but not limited to VLSI, performance modeling and analysis, computer architecture, and influence of computer architecture on operating systems and compiler software.

Our department of nineteen tenure-track and two lecturers emphasizes research, and attracts excellent Ph.D. students, virtually all of whom are fully supported. Departmental facilities include several Sun 4 servers, real-time 3D shaded graphics systems, numerous Sun, HP, Digital, IBM, and NeXT workstations, state-of-the-art equipment for vision and robotics, plus equipment designed and built at Columbia for parallel computation. The department is in the first year of an NSF CISE infrastructure grant and we expect to purchase several parallel processors during the next five years. We are within an hour's drive of the research laboratories of IBM, AT&T, Bellcore, Siemens, Philips, NYNEX, and other leading industrial companies.

Columbia University is one of the oldest universities in the United States, and New York City is one of the cultural, financial, and communications capitals of the world. Columbia has an enclosed campus with grass and tree-lined walks in the Morningside Heights enclave on the Upper West Side. The department is housed in its own building, and in 1992 we will acquire additional space and facilities in the interdisciplinary Center for Engineering and Physical Science Research now in final-stages of construction. University-subsidized housing and parking is readily available.

Candidates for assistant professor should exhibit exceptional research promise, while those seeking a more senior position should have an outstanding record of research achievement. Interest and ability in teaching undergraduates and graduates is necessary. For consideration, we require the submission of a summary of research interests, resume, email address, and the names of at least three references to: Prof. Kathleen McKeown, Faculty Search Chairperson, Department of Computer Science, 450 Computer Science Building, Columbia University, New York, New York 10027 (kathy@cs.columbia.edu).

Columbia University is an Equal Opportunity/Affirmative Action Employer. We encourage applications from women and minorities.

UNIVERSITY OF WISCONSIN-STOUT

Faculty position in computer science available August, 1992. Teach undergraduate courses in computer science; scholarly activity and other faculty duties expected. Ph.D. in computer science or related area required. Relevant experience desired. Salary commensurate with qualifications/experience. Screening begins April 15. Applications close May 4. Send application letter, vita, transcripts and names of three references to: William McGuire, Chair, Math Department, UW-Stout, Menomonie, WI 54751. List of all applicants will be released upon request. EEO/AA.

INFORMATION SCIENTIST

A health information systems provider is recruiting for an information scientist to design, code and test advanced hospital information systems using HELP hospital information systems and knowledge based systems concepts. Design, code and test the MIB for interfacing to bedside monitors and software for interfacing to various systems. MS in Medical Informatics with two years software engineering experience, or MS in Computer Science with four years experience in health care products, 2 years experience in Tal, C languages and Tandem, Charles River computers are required. Experience in design, code and test MIB, knowledge of various standard communication protocols, knowledge of HELP hospital information systems and knowledge based systems are necessary. 40 hours per week, Monday through Friday; $19 per hour. Send resume to:

Job Service
P.O. Box 11750
Salt Lake City, Utah 84147
Job Order #2827859

THE VIRGINIA MILITARY INSTITUTE
Mathematics / Computer Science
Lexington, VA 24450

A tenure-track position beginning August, 1992. Applicant should have a strong interest in teaching.

Preference will be given to applicant with a Ph.D. in computer science. Applicants with significant progress toward a Ph.D. will be considered. Duties include teaching both computer science and mathematics. Salary and rank commensurate with qualifications.

VMI is state-supported with 1300 undergraduates in engineering, liberal arts, and science. It is located in an attractive college town with three colleges within a six mile radius. Faculty wear uniforms but have no other military duties.

Review of applications will begin on March 1, 1992 and continue until position is filled. Send resume, three letters of recommendation, and a graduate transcript to George Piegari, Department of Mathematics and Computer Science, Virginia Military Institute, Lexington, VA 24450.

AA/EEO Employer.

RESEARCH ASSOCIATE

Research Associate need for development of hospital based computer image system. Require Masters Degree in Electrical Engineering or equivalent. Must have knowledge of 1) computer image processing and recognition; 2) physiology, biomedical experiment design, medical decision making; 3) physics and mathematics including optics, discrete mathematics, statistics; 4) immunocytochemical stain theory, PAP-DAB and hematoxylin stain for prostrate, recognition of tumor and normal glands; 5) biomedical imaging, PC286 based light microscopy image system including spectrometry in histology. Demonstratable ability in programming in C, PASCAL, ASSEMBLY; programming the interface of PCVISION image board and PC, BIOS 80286; applying Q-2 image processing, LINC/TCP and PC-NFS workstation software; MS-DOS and UNIX; PC286, SUN, and their communication through Ethernet. Salary $27,000.00 dollars per year. Job Order No. 2828213. Contact Utah Job Service 5735 South Redwood Road, P.O. Box 11750, Salt Lake City, Utah 84147-0750.

GRIFFITH UNIVERSITY, AUSTRALIA
Associate Professor
Software Engineering/
Computer Systems

Griffith University wishes to fill a tenurable Associate Professorship in the School of Computing and Information Technology (CIT), to increase its strength in the SE/Computer Systems area. CIT's major research strengths are in algorithms, databases, programming methodology, software quality, AI, IS, and social aspects of IT. It is responsible for undergraduate, Honours, MPhil, and Ph.D. degrees, and a postgraduate diploma. CIT is a participant in the newly awarded Distributed Systems Technology Centre (DSTC), which has a budget of A$60 million over 7 years. The main research areas of the DSTC are architectures, tools, databases, and performance. The School hosts the Australian Software Quality Research Institute. There are 28 full-time faculty positions, supported by a LAN of over 40 X-terminals and Sun workstations and servers. CIT has access to significant parallel and distributed computing systems. Equipment for teaching includes labs of X-terminals, Macs and PCs.

Applicants should have a Ph.D. an outstanding research record, and teaching experience. They should be active researchers in SE or parallel and distributed systems. The appointee will provide research leadership, contribute to teaching programmes, and assist with administration. He or she will be expected to play a leading role in the DSTC: initiating, directing and contributing to collaborative research projects. The salary scale is A$62525 to A$66625. A market-related salary loading may be offered. Applications should quote Ref. F2-92, include a CV and the names, addresses (including e-mail) and fax nos. of at least 3 referees, and be addressed to the Divisional Administrator, Mr. James R. Walden, Division of Science and Technology, Griffith University, Queensland, Australia 4111. Enquiries should be directed to the Head of School, Prof. Paul Pritchard on tel. +61 7 875 5010 or fax +61 7 875 5198 or e-mail pap@cit.gu.edu.au. The closing date for applications is 30 April 1992. Griffith University is an equal opportunity employer.

SENIOR SYSTEMS ENGINEER

Member of technical staff. Responsible for direct customer support, software design and development, compilers, software development tools, system tests and implementation, develops instructional material as part of design process. Conducts related user training. Must have BA or BS in Engineering and 8 years experience in computer system design and development. Total software experience must include 5 years with personal computers, C language, UNIX based machines, UNIX/AIX and network interfaces and 2 years experience in AIX communications, RISC system/6000. Motif, X Windows. Strong background in IBM Mainframe computers. 40 hours per week, 8:00 a.m. to 5:00 p.m. $96,000 per year. Apply at the Texas Employment Commission, Austin, Texas, or send resume to the Texas Employment Commission, TEC Building, Austin, Texas 78778, J.O. #6687522. Ad Paid by an Equal Employment Opportunity Employer.

RENSSELAER POLYTECHNIC INSTITUTE

The Computer Science Department invites applications for faculty positions at all academic ranks. Postdoctoral and visiting appointments may also be available. Applicants should have a doctorate in Computer Science (or a related area) and a commitment to excellence in teaching and research. Preferred research interests include compiler technology, computer architectures, large scale scientific computing, and programming languages; experience in computer science applications to environmental sciences, interactive learning, or manufacturing process, is of particular relevance. However, all areas will be considered. The Department offers B.S., M.S., and Ph.D. degrees in Computer/Science, and has excellent computing facilities. Send resumes and at least three references to Chairman, Department of Computer Science, Rensselaer Polytechnic Institute, Troy, New York 12180-3590. Rensselaer is an Equal Opportunity/Affirmative Action Employer.

SOFTWARE ENGINEER

Applicant will design and develop software front-end tools to interface with user and company's data base. Will utilize C, SQL languages, database process language, graphic kernel system, and computer graphic interface, and other graphic tools in this process.

Hours and salary: 8:00 a.m. - 5:00 p.m. Mon. - Fri. $36,000/year.

Minimum requirements: 2 years' experience in job offered. M.S. in Computer Science. Must have working knowledge of C, SQL languages, graphic tools (GKS, CGI), relational database systems, operating systems Unix and VMS, "Nested Relation" RDBMS, 4 GL, Pick Basic, Fortran, Pick Operating System, and CAD.

Must have proof of legal authority to work in the U.S.

Attn: Employment Programs
Colorado Department of Labor and Employment
600 Grant Street. Suite 900
Denver, Colorado 80203-3528
Refer to Job Order No.: CO3773908
Please reply no later than 4/30/92.

MIS ANALYST

MIS Analyst needed to design, implement and maintain all insurance related software. Integrate and troubleshoot Local Area Network hardware and software problems. Use Computer Aided Software Engineering (CASE) tools to automate the analysis and design process. Supervise and assist support team members in the installation and testing of special software environments required, such as expert systems, artificial intelligence, computer aided software engineering, decision support systems, and mainframe computing environment being used on the IBM System 36 hardware system. Requires a Bachelor's degree in Computer Science and one year experience in job offered or one year related insurance software experience. Must have one year experience in each of the following: insurance laws and regulations; insurance software programming and local area network. 40 hour work week. $26,890.00 per year. Apply at the Texas Employment Commission, Dallas, Texas, or send resume to the Texas Employment Commisson, TEC Building, Austin, Texas 78778, Job Order #6422249. Ad Paid by an Equal Employment Opportunity Employer.

BOOK REVIEWS

Editor: Alan Kaminsky, Rochester Institute of Technology, PO Box 9887, Rochester, NY 14623-0887
(716) 475-5255; Internet ark@cs.rit.edu;Bitnet arkics@ritvax.

Neural Networks and Fuzzy Systems: A Dynamical Systems Approach to Machine Intelligence

Bart Kosko (Prentice Hall, Englewood Cliffs, N.J., 1992, ISBN 0-13-611435-0, 449 pp., $49)

Neural networks and fuzzy systems are numerical, model-free approaches to the problems of imprecise descriptions and uncertainty in artificial intelligence. They used to be presented as two different subject areas. This book opens a new perspective by examining them from a unified view and demonstrating how they can be combined to construct adaptive systems for AI applications. Specifically, the author shows how to use neural networks to fine-tune fuzzy systems and thus produce adaptive systems.

The book contains two major parts. The first part introduces neural network theory. The material here is broad enough, but the examples tend to illustrate concepts more than applications. The reader may want to look at Kosko's companion book for details on applications: *Neural Networks for Signal Processing* (Prentice Hall, 1992). Nevertheless, the coverage is adequate for handling neural networks in the real world.

The second part is on fuzzy systems, and it's the more remarkable part of the book. It focuses on adaptive systems instead of general fuzzy theory, so readers who want a broader view should consult other references. In addition to introducing the new sets-as-points theory for interpreting fuzziness, the author demonstrates how adaptive fuzzy systems work by presenting three applications: backing up a truck, image transform coding, and target-tracking controllers. Actually, these three applications respectively benchmark the adaptive fuzzy systems against three popular approaches to dynamic control systems: pure neural networks, dedicated algorithms, and Kalman filter (stochastic math-model).

The author states that his basic theory uses only material from upper-level undergraduate curricula in engineering and science; however, in my opinion, depth and fluency in this material may be required for the reader to fully appreciate this book. The presentation is sometimes concise in the way of technical reports or journal papers. Readers new to neural networks or fuzzy systems or who are used to self-instruction textbooks may find some parts too dense with definitions.

The exercises at the end of each chapter are well organized and cover a reasonable range of difficulty. Some exercises require coding a simple program to conduct an analysis, and some require running the software on the accompanying floppy disks. This is indeed the best way to lead the reader to in-depth comprehension in such computer science subjects. The software requires an EGA card, which might be a limitation for some users.

Due to the relative novelty of combining neural networks and fuzzy systems, I would recommend this book for upper-level undergraduate or graduate courses on special topics or as a complementary textbook for courses dedicated to either neural networks or fuzzy systems. The book is also broad enough and sufficiently self-contained to serve as a reference work for computer professionals interested in these areas.

Tsai-Chia Peter Chou
Cambridge Research Associates
McLean, Virginia

Microcomputer Buses

R.M. Cram (Academic Press, San Diego, Calif., 1991, ISBN 0-12-196155-9, 246 pp., $49.95)

The author intends this book for system and design engineers. He also claims that engineering students will find it useful as an "introduction to the basic features and capabilities of a range of popular commercial buses."

The book is structured accordingly. Three introductory chapters cover most basic concepts and background information for prospective bus users. Buses are compared mostly in terms of protocol characteristics (data transfer and interrupts, arbitration, and remaining functions like error handling). The discussion is interesting, but I know at least one writer who has been more informative on these characteristics (see Paul Borrill's articles in *IEEE Micro*, Vol. 5, No. 7, Dec. 1985, pp. 71-79, and in *Microprocessors and Microsystems*, Vol. 10, No. 2, Mar. 1986, pp. 94-100).

The introduction includes a valuable discussion of hardware and software interfacing. I agree with the author that it is difficult to find "the magic individual who is capable of bridging the gap between hardware and software," but I was disappointed that so few pages were dedicated to this. I would like to see the hardware/software section significantly expanded, just because there are very few real experts in this field. Probably due to a somewhat artificial separation between electrical and computer engineering departments and computer science departments, the professional world is divided into hardware engineers and software engineers who may not have the best communication with each other. Hardware and software interfacing will continue to be a crucial problem for system integration. The book would have benefited from more interfacing examples — say, in mixed assembly, C, and Ada — like those given for assembly and Pascal.

The core of the book is a review of six buses: four that are used predominantly in industrial environments (Multibus I and II, VMEbus, and STD bus) and two that are used predominantly in personal computer markets (AT bus and NuBus). The bus descriptions concentrate on mechanical and electrical specifications, followed by an outline of signal lines, their functions, and the bus protocol characteristics discussed in the introduction. The author deliberately avoids detailed descriptions. In some places the descriptions are even incomplete in that they do not follow the original bus specifications. He explains that "the intent of this book is not to provide the same level of detail present in the national specifications but instead to offer a simplified view of each bus and its most important characteristics." Nevertheless, I felt that some of the more important concepts — for example, block (sequential) operations for Multibus II — should be treated more extensively. Also, engineers familiar with buses will discover with sorrow that Micro Channel and Futurebus descriptions are missing.

The book concludes with four chapters that address design issues, including board interfacing and engineering. I thought the last part of the book was really where the author should have focused his effort. Its first two chapters describe programmable logic arrays and gate arrays in the context of board interfacing to the bus. The last two chapters present simple but instructive board design examples, one for Multibus I and one for VMEbus. Unfortunately, they look like peripheral subjects left to the very end. I think these should have been central topics. Bus descriptions themselves can never compete with the hundreds of pages of detailed specifications. They should form only a background, and system integration topics should get top priority.

Only two of the topics that become real problems in bus design are discussed here: transmission line effects and metastability. There are many more. True, the author does not intend to serve bus designers' needs. Nevertheless, board designers and system integrators need to be aware of such topics as transceiver technology, skews, bus locking and preemption, and cache coherence. A more serious flaw is the complete absence of references. This is simply a mistake.

In the context of other books on this subject, I would categorize this one as very introductory — more valuable to students than to professional engineers. It cannot compete with two comprehensive texts on buses: D. Del Corso et al., *Microcomputer Buses and Links* (Academic Press, 1986) and *Digital Bus Handbook*, J. Di Giacomo, ed. (McGraw-Hill, 1990). These texts, however, are addressed mostly to bus designers. For this reason, I believe Cram's book should have concentrated even more on the viewpoints of system integrators and board designers. This would have meant further enhancement of the chapters on circuit design, board design, and hardware/software interfacing, and moving bus descriptions to the background. Nevertheless, Cram presents a good introduction to an area covered by no other book on the market today.

Janusz Zalewski
Southwest Texas State University
San Marcos, Texas

High-Performance Computer Architecture, Second Edition

Harold S. Stone (Addison-Wesley, Reading, Mass., 1990, ISBN 0-201-51377-3, 459 pp., $53.25)

Several important issues are central to today's high-performance computer architectures, and this textbook covers most of them. It is an extremely well written classic text, suitable for advanced undergraduate and first-year graduate students. It assumes a background in basic computer organization, elementary probability theory, and programming in a procedure-oriented language (for example, Pascal).

Memory system design is the first important consideration in computer architecture. The text gives a good explanation of program characteristics exploitation (temporal and spatial localities), then discusses cache and virtual memories. Performance analysis and modeling for caches under various work-load conditions can be done by either trace-driven simulations or mathematical models. The text gives a good account of these topics including address trace generation, trace stripping, and transient effects. It also briefly describes multilevel caches and discusses several issues regarding virtual-memory structure, mapping, and replacement policies.

The chapter on pipeline design examines, in depth, instruction pipeline design, memory structures supporting pipeline operations, and control of pipeline stages in a multiple-function pipeline. It presents many good examples from existing architectures, but the discussion of pipelined reduced instruction-set computers is barely minimal.

Computer architects must understand large-scale numerical applications and their mapping to particular architectures. Physical computational models fall mostly into two categories: a continuum model, which is based on near-neighbor interactions, and a particle model, which is based on discrete point-to-point interactions. Architectures and algorithms for both models are discussed in reasonable detail. There is also a good discussion on vector computers and the mapping strategies of some numerical algorithms.

Two well-written chapters on multiprocessor architectures and algorithms cover interconnections, performance issues, cache coherence, parallel synchronization, and the transformation of serial algorithms into parallel algorithms. Multilevel cache-coherence, a problem that is currently attracting research interest, is an important omission, however. The discussion of SIMD array processors is very limited. Moreover, dataflow computer architectures are not covered at all.

An excellent glossary covers commonly used terms, and the problems at the end of each chapter are carefully written. Some may boggle the minds of graduate students.

The book's minor deficiencies can be made up by other texts such as K. Hwang and F.A. Briggs, *Computer Architecture and Parallel Processing*, McGraw-Hill, 1984). For the most part, this book is full of ideas and information, and offers a real contribution in the computer architecture field. I highly recommend it for students as well as practicing computer engineers.

Rana Ejaz Ahmed
Anaheim, California

Update

The March review of *Software Engineering with Abstractions*, a new textbook by Valdis Berzins and Luqi (Addison-Wesley, 1991), reported that software tools for the Spec language (used in the book's examples) were available from the Open Software Foundation. The author has informed us that these arrangements changed after the book was printed. The Spec tools are instead available via anonymous file transfer protocol from taurus.cs.nps.navy.mil in the Spectools directory.

"Any clod can have the facts, but having opinions is an art."
Charles McCabe, San Francisco Chronicle

THE OPEN CHANNEL

Editor: Will Tracz, IBM, Federal Sector Division, Mail Drop 0210, Owego, NY 13827
e-mail tracz@owgvm0.vnet.ibm.com

New open channels, old-fashioned societies

The Usenet public electronic bulletin board is the real "open channel." Not only is it more "open" than the column you're reading, it is livelier and faster, and carries a lot more data. For years, the cliche was that we don't find out where our fields are headed by reading journals or conference proceedings; we find out by talking to our friends and bumping into people at conferences. Now, we also have Usenet.

No other public bulletin board comes close to Usenet in breadth and bandwidth. It has nearly 2 million readers and carries about 25 megabytes per day, and it is still growing exponentially. Although it uses traditional international networks like Internet and Bitnet for data transport, Usenet itself has no central organization. A few sites sell news feeds and for-profit data, but Usenet mostly consists of volunteers exchanging free data. In particular, I have nothing to gain financially by mentioning Usenet; like most Usenet sites, my company is just a leaf node, and we subscribe mostly to read juicy bits of technical information that we filter automatically from Usenet's oceans of noise.

It's fun witnessing the creation of a new medium. As with the birth of cinema around the turn of the century, there's a heady mixture of genuine innovation, life imitating art, unwanted advertisements, new forms of pornography, cries of alarm from outsiders, confusion about copyright law, unrewarded genius, much wasted time, and periodic predictions of the imminent death of the medium. Most people haven't even heard of Usenet yet, so it has plenty of room to grow; but nobody really knows where it's headed.

You can read Usenet articles by cadging an account from a nearby university or company, by subscribing to a CD-ROM clipping service, or by investing in a home computer, a good modem, and a news feed. If you can't find a feed from a local volunteer, for a few bucks you can subscribe to UUnet (Falls Church, Virginia), a nonprofit organization devoted to Usenet, or to any of several profit-oriented news-feed companies.

Traditional organizations are slowly becoming aware of the electronic bulletin board phenomenon in general and Usenet in particular. The news media rarely report on it, perhaps because they are wary of the potential competition. But what I find more interesting is the stance of professional computer groups like the IEEE Computer Society and the ACM.

At first, you might think that computer societies would eagerly embrace the use of computers to broaden communication. But their track record suggests otherwise. Electronic mail was used for many years before e-mail addresses started creeping into computer journal mastheads. Even today, it's a rare computer workshop that makes its proceedings available electronically via computer networks. So, it's not surprising that the computer societies mostly ignore Usenet. The reasons relate to money and inertia.

Money is what makes professional societies run, and the common perception is that they can't make money off Usenet. That's too bad, because Usenet could save them money. For example, they could distribute election material and newsletters via Usenet to members like me who prefer to receive them that way; this would save printing and postage costs. They could also distribute money-making publications like journals and standards via commercial Usenet services, which already carry material ranging from market surveys to syndicated humor columns.

I'd certainly prefer storing my professional journals on a few small gigabyte tapes or disks and shrinking my 10-meter bookshelf of paper journals, which are constantly being misplaced and are hard to search through.

Inertia is a more fundamental problem. The current generation of professional society movers and shakers grew up in a paper-oriented world, and most of them don't know Usenet at all or view it with curmudgeonly disdain. As a whole, they seem to lack motivation to use this new electronic medium that computing helped create. But if they don't wake up soon, they may get left behind. New societies form when old ones get lazy.

Paul Eggert, Twin Sun
El Segundo, California
eggert@twinsun.com

Computer editor's note:

The IEEE Computer Society is definitely interested in using alternative mechanisms to distribute technical materials and is actively involved in various efforts. While we all agree that the Usenet public electronic bulletin board is an extremely valuable and timely means of discussing issues, the Computer Society has felt that it is insufficient as a means for communicating to the membership at large.

The Computer Society's first requirement was to ensure that all members had available basic interconnection services. To accomplish this, the society established the Compmail service — complete with Internet access. While electronic mail was the initial service provided by Compmail, the long-term goals are value-added services such as on-line conference information and registration service, and advance abstracts.

The Computer Society is also interacting with the standards working groups to assist in using electronic networks to speed the development of draft standards. In addition, the society is working with the IEEE to establish an IEEE electronic database that will include the above services and be accessible through Internet and Compmail. The key to all services that the Computer Society offers is guaranteed and easy access for all members.

The real issues that have slowed the introduction of on-line journals are not technical but economic. The society must be able to recoup costs associated with any value-added service. Ultimately, do users want the benefits inherent in electronic distribution and are they willing to pay for it? All society efforts must be directed at answering these questions.

Anyone interested in participating in these efforts should contact James H. Aylor, Computer Society president-elect and chair of the Electronic Media Development Committee; e-mail j.aylor@compmail.com; phone (804) 924-6100.

COMMNEXUS

Presents

The LiSBUS™ Async I/O System

A solution of the 1990's for today's data transmission problems.

Outstandingly Simple and Reliable because **LiSBUS**™ is based on a breakthrough technology which uses the impedance of the bus cable to replace binary addresses. Consequently, data transmission management is greatly simplified and much more reliable than today's equivalent systems which require expensive software, hardware, and personnel investments.

Outstandingly Practical because it is easy to install and operate. No special tools, workbench, or electronics expertise are needed. Anyone can be up and running in minutes. Just plug in the external modules and configure the system with the user-friendly **LiSBUS**™ **Link Control Software.** Each external module measures only around 2in. by 2in.

Outstandingly Flexible because a user can connect up to 60 peripherals or computers to a controlling computer through their RS-232C (COM) ports. To add peripherals, just extend the bus cable and add modules.

At an Unbeatable Price because at $650* for the **LiSBUS**™ **Starter Pack**, no alternative offers all these advantages combined into one product without spending a much higher amount. The **Starter Pack** includes all the user needs to connect four peripherals and a complete set of **LiSBUS**™ **Software Development Tools** to create custom applications.

LiSBUS™ **Async I/O System:** A product of our **CommNexus**™ line of communication systems. GIGATEC is committed to offering products and servicing its customers in the best tradition of Swiss quality. We provide our customers with:

- **Technical Support.** Registered buyers can obtain technical support from our qualified engineers.
- **Users and Developers Group.** Organized for encouraging software developments using the products of the **CommNexus**™ family.

For more information and ordering contact:

In the USA and Canada:

Toll Free (800) 945-3002
(excl. Hawaii)
Mon.-Fri. 9am - 9pm EST

GIGATEC (USA). Inc.
871 Islington Street
P.O. Box 4705
Portsmouth, NH 03802-4705 USA
Tel. (603) 433-2227
Fax (603) 433-5552

In Europe:

GIGATEC SA
Ch. des Plans-Praz
1337 Vallorbe SWITZERLAND
Tel. 41 21 843 37 36
Fax 41 21 843 33 25

* specifications and prices subject to change without prior notification

Visa and MasterCard/EuroCard accepted. - CommNexus™ and LiSBUS™ are trademarks of GIGATEC SA.

Reader Service Number 4

Department of Computer Science

Stanford University

Ten intensive computer short courses presented by an outstanding faculty. The courses are designed to give attendees an immediately useful increment of knowledge and skill.

For information, call (916) 873-0575

July 20-24
Advanced Geometrical Computing
Leonidas Guibas, Daniel Huttenlocher, Jean-Claude Latombe

July 20-24
Object-Oriented Methodology
Kenneth S. Rubin, Vicki L. Katzman

July 27-31
Advanced Database Technologies: Object Data Management
Arthur Keller, Rick Cattell, David Maier, others

July 27-31
Computer-Aided Software Design, Development, and Maintenance
Ellis Horowitz, Anthony Wasserman

July 27-31
Compiler Construction/Programming Language Translation
Susan L. Graham, John L. Hennessy, Jeffrey D. Ullman

August 3-7
Code Optimization and Code Generation
Susan L. Graham, John L. Hennessy, Jeffrey D. Ullman

August 3-7
Introduction to C++
Andrew Koenig, Barbara Moo

August 3-7
Telecommunications for Computer Professionals
Dale A. Harris

August 10-14
Parallel Computer Architectures and Compilers
Anoop Gupta, Monica Lam

August 17-21
Recent Trends in Processor Design: Reclimbing the Complexity Curve
Mark Horowitz, Steven Przybylski, Michael D. Smith

WICS

FOR FURTHER INFORMATION ON THESE COURSES, WRITE OR PHONE:
Joleen Barnhill, Program Manager
Western Institute in Computer Science (WICS)
Post Office Box 1238, Magalia, CA 95954-1238, (916) 873-0575

Reader Service Number 5